Annual Review of Communications

International Engineering Consortium

Volume 55

The *Annual Review of Communications, Volume 55* includes papers from educational forums sponsored by the International Engineering Consortium, a nonprofit organization affiliated with more than 70 universities worldwide.

Quality Publications from the International Engineering Consortium:

- IP and the Intelligent Network: The Road to Convergence
- Next-Generation Optical Networks: Advanced Technologies and Market Opportunities
- Operations Support Systems 2002: Enabling the Next-Generation Network
- The Evolution of Broadband: DSL and Beyond
- Optical Networks: Architectures and Advances
- Optical Access: Next-Generation Technology and Services in the First Mile
- Telecom Outlook Report: Millennial Edition
- DSL's Next Chapter: Packaging Access, Adding Voice and Video
- The Market for IP Enhanced Services
- Broadband Access and Services in the Local Loop: A Comprehensive Report

For more information on any of these titles, please contact the IEC publications department at +1-312-559-3730 (phone), +1-312-559-4111 (fax), publications@iec.org, or via our Web site (http://www.iec.org).

This *Annual Review of Communications,* together with subsequent editions, may be purchased from the International Engineering Consortium.

Volume 1, 1944, through Volume 43, 1989, are out of print. Microfilm copies of these volumes may be obtained from University Microfilms, Inc., Ann Arbor, Michigan 48106.

ISBN 0-931695-02-4

International Engineering Consortium
549 West Randolph, Suite 600
Chicago, Illinois 60661-2208 USA
+1-312-559-3730 phone
+1-312-559-4111 fax

Contents

Section II: Business, Marketing, and Regulatory Issues

Section IV: Packet Networking, IP Telephony, and the Internet

Section V: Wireless

Contents by Author

About the International Engineering Consortium

The International Engineering Consortium (IEC) is a nonprofit organization dedicated to catalyzing positive change in the information industry and its university communities. Since 1944, the Consortium has provided high-quality educational opportunities for industry professionals, academics, and students.

In conjunction with the industry, the IEC has developed free, on-line, Web-based tutorials. The IEC conducts industry-university programs that have substantial impact on curricula. Every year, the IEC provides grants to professors and students to participate in the IEC educational forums. The IEC also conducts research and develops publications, conferences, and technological exhibits that address major opportunities and challenges of the information age.

More than 70 leading, high-technology universities are currently affiliated with the Consortium. The industry is represented through substantial corporate support and through the involvement of many thousands of executives, managers, and professionals.

University Program Sponsors

Through the University Program, sponsoring organizations provide grants for full-time faculty members and their students to attend IEC Forums. The generous contributions of the following sponsoring organizations make this valuable program possible. For more information about the program, please call +1-312-559-4103 or send an e-mail to *universityprogram@iec.org*.

This past year, the University Program has provided more than 400 grants to professors to attend the educational opportunities offered at IEC Forums. Because of their participation in IEC Forums, professors created or updated 600 university courses; more than 10,000 students were impacted by these new and upgraded courses, and more than 12,000 students benefited from improvements made in university laboratories. Since its inception in 1984, the University Program has enhanced the education of more than 500,000 students worldwide.

Agilent Technologies
cingular WIRELESS
IBM
PACIFIC BELL SBC global network
tellabs

European Sponsor

BT
British Telecommunications plc.

ALCATEL
Cisco Systems
invensys
SBC
TERADYNE

Ameritech
DSL FORUM
JOHNSON CONTROLS
ARTHUR J. SCHMITT FOUNDATION
TRENDIUM

ANDREW
EDS
Lucent Technologies Bell Labs Innovations
SIEMENS
verizon

ARRIS
ERICSSON
Marconi
SNET SBC global network
WORLDCOM

BELLSOUTH
HEWLETT PACKARD
MOTOROLA
TDS TELECOM

Consortium Affiliate Universities

The International Engineering Consortium is affiliiate with the following leading hi-tech universities:

- The University of Arizona
- Arizona State University
- Auburn University
- University of California at Berkeley
- University of California, Davis
- University of California, Santa Barbara
- Carnegie Mellon University
- Case Western Reserve University
- Clemson University
- The University of Colorado at Boulder
- Columbia University
- Cornell University
- Drexel University
- The University of Florida
- Georgia Institute of Technology
- Howard University
- Illinois Institute of Technology
- University of Illinois at Chicago
- University of Illinois at Urbana/Champaign
- Iowa State University
- The University of Kansas
- The University of Kentucky
- Lehigh University
- Marquette University
- The University of Maryland at College Park
- Massachusetts Institute of Technology
- The University of Massachusetts
- Michigan State University
- The University of Michigan
- Mississippi State University
- The University of Minnesota
- The University of Missouri-Columbia
- The University of Missouri-Rolla
- North Carolina State University at Raleigh
- Northwestern University
- University of Notre Dame
- The Ohio State University
- Oklahoma State University
- The University of Oklahoma
- Oregon State University
- The Pennsylvania State University
- The University of Pennsylvania
- The University of Pittsburgh
- Polytechnic University
- Purdue University
- Rensselar Polytechnic University
- University of Southern California
- Stanford University
- Syracuse Univesity
- University of Tennessee, Knoxville
- Texas A&M University
- The University of Texas at Austin
- VA Polynechnic Institute and State University
- The University of Virginia
- The University of Washington
- The University of Wisconsin-Madison
- Worcester Polytechnic Institute

International Affiliate Universities

- École Nationale Supérieure des Télécommunications de Bretagne
- École Nationale Supérieure des Télécommunications de Paris
- École Supérieure d'Électricité
- Imperial College of Science, Technology and Medicine
- Institut National Polytechnique de Grenoble
- Instituto Tecnológico y de Estudios Superiores de Monterrey
- KAIST
- McGill University
- The Queen's University of Belfast
- Technische Universität München
- Universidad Nacional Autónoma de México
- Université d'Ottawa
- University College London
- University of Edinburgh
- University of Glasgow
- University of Southampton
- University of Toronto

Media Sponsors and Partners

The IEC's Media Sponsors and Partners program was established as a means for the Consortium and the leading industry publications to provide reciprocal support for the many service programs throughout the year. The IEC thanks these publications for their efforts in keeping the industry up to date on the Consortium's events, publications, and Web-based educational programs.

Media Sponsors

Publications participating as Media Sponsors provide substantial support of and exposure to the Consortium's overall programs and are intimately involved in specific events and publications as Official Publishing Sponsors.

America's Network
201 East Sandpointe Avenue, Suite 600
Santa Ana, CA 92707-8700
714-513-8400
www.americasnetwork.com

Communications News — Solutions for Today's Networking Decision Makers
2500 Tamiami Trail North
Nokomis, FL 34275
941-966-9521
www.comnews.com

Telephony
One IBM Plaza, Suite 2300
Chicago, IL 60611
312-595-1080
www.internettelephony.com

boardwatch — Technology and Business for Communications Service Providers
13949 W. Colfax Avenue, Suite 250
Golden, CO 80401
1-800-933-6038
www.ispworld.com

LIGHTWAVE
98 Spit Brook Road
Nashua, NH 03062-5737
603-891-0123
www.light-wave.com

Wireless Week
8878 South Barrons Boulevard
Highlands Ranch, CO 80129-2345
303-470-4875
www.wirelessweek.com

RCR Wireless News
777 East Speer Boulevard
Denver, CO 80203-4214
303-733-2500
www.rcrnews.com

Media Partners

Publications participating as Media Partners provide general support to the Consortium's overall programs.

Broadband Wireless Business Magazine
www.shorecliffcommunications.com/magazine/news.asp

The Net Economy
www.theneteconomy.com

telecomasia
www.telecomasia.net

Business Communications Review
www.bcr.com

Outside Plant
www.ospmag.com

Telecommunications Magazine
www.telecommagazine.com

commverge
www.commvergemag.com

PHONE+
www.phoneplusmag.com

Global Telephony
www.globaltelephony.com

Private & Wireless Broadband Magazine
www.privatebroadband.com

Section I:

Executive Perspectives

The Light Bulb, the PC, and Broadband

Duane Ackerman
Chairman and Chief Executive Officer
BellSouth Corporation

As recently as two years ago, not many folks were interested in what communications companies from the "old economy" had to say about the Internet. Nor did they want to hear our take on the digital revolution more generally. After all, we were old economy. All we had were customers and revenue, strong balance sheets, big networks, billing operations, and a long heritage of service.

And it didn't seem to matter that, for a more than a century, we had been taking the latest in communications technology, and turning it to the advantage of customers—while also delivering earnings to our shareholders. As a business writer for "The New York Times" said in the spring of 2001, "old economy measures looked hopelessly quaint" back then.

That was then, this is now. Today, traditional measures such as earnings are back. Likewise, it is understood once more that the ultimate discipline is imposed by the customer—the products and services they'll use, the price they'll pay, the level of service they demand. And folks think companies such as BellSouth might actually be the ones to deliver on the digital revolution. But they're not sure.

They ask: Who's got a business model that will work? They know the race is on to deploy broadband—"fast access." And they recognize the growing consensus on what will ultimately determine the success of broadband: "It's the applications, stupid."

And while they know BellSouth and others are committed to enabling the broadband network, they ask: Where are the applications? Can anyone make money on them? Or is all of this just more hype, like so much hype we've heard for more than 10 years, ever since talk about the "information highway" became the rage?

Good questions all. I can answer them emphatically: No, broadband and the rest are not all hype. Someone does know how to separate the practical from the hype, the profitable from the hype.

Even in an economic slump, we see the promise of substantial growth ahead of the industry. Customers are adapting a new generation of digital appliances and services, and in typical digital fashion: Costs are going down; performance is going up. Convergence is changing network economics, creating a new competitive dynamic.

The pending arrival of broadband/Internet protocol (IP) infrastructure, through digital network convergence, coupled with applications, will drive increased customer demand. We see this supply-side effect as a potent driver of demand, including demand for full packages of service. And just as the personal computer (PC) ushered in a new era, so too will broadband—a new era for the economy, a new era for the industry. The fact is we have reached the limits of narrowband Net growth—broadband" will spur the next wave of growth.

The Applications Story

The PC is a powerful example of what will happen with broadband applications. For many of us, the PC started as a glorified typewriter—a word processor. Applications followed. Applications followed when the PC was a standalone device. Applications followed when it was used with local-area networks (LANs) and wide-area networks (WANs), and more applications came with the Net. They're still coming. The applications story will continue in that vein with broadband.

The story has been the same for other technologies: television, radio, and electricity applications—successful mass applications. Of course, there were always skeptics. Darryl Zanuck, the head of 20th Century Fox, said this about television back in 1946: "Television won't be able to hold on to any market it captures after the first six months. People will soon get tired of staring at a plywood box every night."

And back in 1977, a computer executive—who I won't name—said this: "There is no reason for any individual to have a computer in their home." I suppose he was right—in 1977. But we know that the increase in the power of computers, the decline in price—and the development of applications—made a shambles of that 1977 projection.

As we see the world, the skeptics on broadband will be just as wrong. We at BellSouth see substantial promise for growth in broadband, in communications generally.

Now, as you here know, the prevailing mood in the industry isn't optimistic; it sure doesn't feel like we're an industry with growth on the horizon. We're all watching that stock ticker more closely, following the Fed and Alan Greenspan more closely, looking at manufacturing stats, sales of durable goods, new housing starts, unemployment—all those tea leaves. We know the market is edgy. But as we look at market demand and as we look at the prospects for broadband, we have reason to be more optimistic than today's glum mood indicates. Overall, data growth is running at 30 to 40 percent. In BellSouth's case, with some specific data products, we're seeing 50, 60, 70 percent growth on what are now significant bases. The same is true for other companies.

Digital Story Still a Good Story

People like the power and flexibility that digital devices allow them. They're using them; they want more.

The nub of the matter is this: Even as the short term concerns everyone, the digital story looks good for the long term.

In the emerging broadband/IP world, different parts of the industry will have new and somewhat differing advantages as we pursue high-value customers of network services. These advantages will differ especially depending on the specific assets with which each of us starts. And hence our digital stories will differ.

The BellSouth story includes some big numbers that touch all of you as well—numbers on growth in the total volume of information. These give us a better sense of what's happening in the big picture of data. These numbers come from the School of Information Management and Systems at the University of California at Berkeley. And while we now need to view these numbers in light of this economic slump, they nevertheless suggest the scale of data growth that we need to be thinking about.

The school estimates that the world's total information at close of 2000 was roughly 15 billion gigabytes. By the end of 2002, the school projects total information will be 33 billion gigabytes. In other words, in just two years, there'll be a doubling of the amount of information the human race has accumulated since humans first daubed paint on cave walls. By 2003, the amount will stand at 57 billion gigabytes.

No wonder digital customers are clamoring for data products.

All of this information has to be generated—perhaps processed, manipulated some way, maybe with other information in other media. It has to be transported, stored, retrieved. In other words, not only is there a lot more data in our future, but more of it is being moved more often and at ever-higher speeds.

That's a lot of work for everyone in the information-communications industries—a lot of applications to develop.

With this continuing explosion of information, storage and storage-related solutions are projected to be a near $100 billion market within three years. Businesses need new efficiencies, new ways to increase productivity. Consumers always want more convenience, new ways to gain control of their time. Give them broadband applications that deliver on these—applications that are practical, powerful and affordable—and they'll use them.

Practical, Profitable Broadband

The following are several areas we at BellSouth believe are the most promising, the most practical, the most profitable.

The first is what we call "productivity-enhancing" applications for small and mid-size businesses: enterprise resource planning (ERP), customer-relationship management (CRM), and front- and back-office applications.

Our long economic boom has been driven by these applications for large businesses. Now, we can address medium to small businesses. Such applications would be in areas where small and mid-size businesses simply cannot afford the information technology (IT) infrastructure to do the work, but where they want some of the same services available to big companies. By 2004, this is expected to be a $15 billion market in the United States, a third of that in small and mid-size businesses.

Video is another promising area. Training for corporations is of course a "videoconferencing" kind of opportunity. Training falls under the label of "distance learning." Originally, we thought of distance learning as being for K-12 education as well as some higher education. As it turns out, 75 percent of the distance-learning market is corporate training. By 2004, this market will be around $30 billion in the United States.

We all recognize the growing challenge that training presents to us right now—new people plus new technology. Companies are increasingly turning to a video solution for training—especially when, as is so often the case, employees are dispersed in many locations.

And there are practical video applications for consumers. "Instant" videoconferencing can be used for security. From the office you might use video to check things at home or to see if the delivery driver left the package in the right place. Similarly, as your parents age, you can use video to check with them throughout the day. Likewise, if you have a weekend house, a boat stored at a lake, or a new babysitter that you're not yet sure of.

We see instant video messaging as a promising area for broadband, for instance through sharing files with family members. Grandparents get the latest videotape of the grand babies, the tape of the school play, and so on.

Video programs on demand are another area, as well as interactive video games. On-line games alone are expected to be a $10 billion market by 2004. And we're looking at various devices for the home in which an assortment of movies and other programs might be downloaded overnight, rather than playing in real time. We believe this is a promising interim solution.

We expect customers to take to unified messaging (UM): Bringing together all the devices, all the networks, so someone can sit down at their PC and listen to their voice-mail—with one mailbox for wireline and wireless—as well as read their e-mail. Or someone can use their wireless phone or Palm Pilot, or other devices to retrieve other messages and information.

And then there's growth in what we're all calling the "enablement layer," that being the additional capabilities the network requires. To offer video on demand (VOD), for example, takes this enablement layer. We have to have distribution networks with the broadband capacity to transmit video. And we have to have storage to cache movies—or other content—near the edge of the network. At BellSouth, we see this enablement layer as a major opportunity. We're making such "enablement" an integral part of our network as we transition our network from circuit to packet, from analog to digital, from narrowband to broadband.

And here's another bright spot for broadband: Reliable, secure broadband access to all the existing applications that reside on enterprise networks—a "remote access" VPN.

These are just a few examples of some specific, practical broadband applications that we believe are promising.

Such applications, of course, beg the question of who profits from them—in particular, how will network companies benefit? Can we construct a win/win solution? At BellSouth, we're committed to various kinds of alliances and partnerships in this regard. In fact, we're working with many on just such matters.

Mass Technologies and Applications

We can all take additional encouragement from the long experience already alluded to, the history of earlier technologies. Take the development of the PC from standalone to now, when, through the Internet, it can be hooked to other computers and to databases around the world. Applications were key; the PC became not only powerful, but also practical.

Applications drove success in other mass technologies.

Anyone of a certain age can remember when television went off at 10 o'clock in the evening. After that, all you could get was a test pattern. In Meridian, Mississippi, the local television station filled much of the airtime with a man sitting there playing the organ live. Look at the content now.

Or electricity. It's a useful analogy. Initially, the primary "application," if you will, was the light bulb. Simple as that was, it revolutionized how people lived. Folks such as Edison, of course, knew there was far more in electricity than light. And indeed, there came a long line of "applications": refrigerators, record players, pumps, stoves, washers, air conditioners—"applications" of electricity that reduced the amount of sheer drudgery in work, improved health, and delivered comfort and convenience. The light bulb was just the first step.

Today, people see broadband as fast access to the Internet; that's all. But fast access is just the first step. We can't foresee all of the transforming types of applications that will emerge from broadband—the practical, multimedia applications. Those will come. We'll see hundreds of them—some of which will change work and leisure as much as did electric pumps, refrigerators, air conditioners, television, and all the rest.

I have faith in that assertion. It is faith built on experience. It is not "build it, and they will come." It is "Make it practical, make it powerful, make it affordable—and they will pay." Real value generates real revenue.

The sum total of our experience tells us that this is the case.

A New Competitive Dynamic

Convergence will deliver much more powerful products and services. Our own view is that convergence, as it moves us toward integrated broadband/IP networks, is also producing a new competitive dynamic. We also believe that convergence remains the most powerful force at work in our industries. Consider wireless, and, more specifically, the fate of local and long-distance service in the wireless arena. We all know what happened. With AT&T's One Rate Plan as the precipitating event, local and long distance merged. They converged.

AT&T radically altered the landscape. Everyone had to adjust: rate plans, bundles, and joint ventures such as Cingular.

Another measure of the power of convergence can be seen in the recording industry. Not long ago, only big companies could distribute music worldwide. Today, a 10-year old kid in, say, Dothan, Alabama, can be a worldwide distributor—much to the chagrin of recording artists and the recording industry.

Convergence is not a single destination. It is a journey—a complex journey, a continuum. Its way stations are network technology platforms, protocols, and services. As we pass the various way stations, each of them interacts with the others, creating still more powerful convergence—and more powerful services for customers.

It's this broader force of convergence that I believe is the key to future growth. Convergence is what makes more powerful packages of services less expensive and, therefore, attractive to customers.

Contrary to much of the current conventional wisdom, we at BellSouth believe that "the package" is alive and well. But we believe that the package is alive and well for a different reason than made the idea so popular during the past few years. Then the idea was that the convenience of "one-stop shopping" would prompt consumers to buy full packages. The new competitive dynamic being created by convergence will give consumers a more potent incentive for buying full packages: lower costs. In other words, the changing supply-side dynamic, not so much demand-side, is the real driver of demand for packages.

We see tremendous opportunities here. There is a place for a full package of services, a big place. These will be highly successful with many segments of customers, particularly those who use the most services. Yes, there'll also be places for players that are not providing full packages. Some will be major players, and some will be niche players. But in the world of integrated broadband/IP networks, we see the offerings of full-service providers becoming more attractive to an increasing range of customers.

And we believe we're positioned to be a full-service provider; our local network, we believe, will be at the heart of delivering integrated, broadband/IP capabilities to consumers and businesses. In BellSouth's case, our existing local network makes us a primary enabler in the digital revolution.

The Digital Challenges

The challenges for our related industries are clear: We know we have to deliver an integrated broadband/IP network. We know we have to deliver the new architectures, which include storage and caching facilities at the network's edge. We know we have to deliver facilities for Web hosting. We know we have to deliver carrier-quality reliability and security. We know we have to deliver applications that are practical, powerful, and affordable. We have to help customers be more productive. We have to find business models that work. We have to construct win/win relationships with our partners. And we have to help the digital customer with the complexity that we've pushed to the edge of the network.

To deliver on all of these requires all of those traditional or "old economy" virtues that we talked about at the start—the knowledge and skills that we've developed over the years and our ability to deliver to customers and to shareholders.

The digital story is a great story, and we're in a great business. When the world is creating a digital economy, an increasingly networked economy, it's pretty good to be in our business. I always tell our employees that we're in the most exciting industry in the world. We're an enabler of the most basic forces driving the economy and driving changes in society—an enabler of forces of historic power and reach.

The opportunities are immense. Our challenge is to execute and deliver, and I believe we can. When we do, we will have earned our spot in history.

This paper was adapted from a speech delivered to the Accenture Global Communications Forum in Miami, Florida on April 2, 2001.

Architecture for Multifunction Application Server

Shankar Allimatti
Corporate Head, Centre of Excellence–Intelligent Networks
Mahindra British Telecom Limited

Introduction

The revolutionary change and the growth in the telecommunication industry is witnessed today in the quest for achieving universal communication—anywhere, any time, anyone. This exponential change phenomenon is basically driven by the new emerging technologies, increased customer demands for better services, global competition across the industry, and time to market leading to mergers and acquisitions (M&A). Above all, the economic compulsions are driving network/service providers and equipment suppliers to increase operational efficiencies and cut down on costs.

With such competitive compulsions and the fact that the existing networks with very large investments are here to stay for some time to come, the network/service providers are aptly choosing the path of orderly migration. This orderly migration means the evolution of existing public switched telephone networks (PSTNs) with a time division multiplexing (TDM)/time division multiple access (TDMA)–based infrastructure optimized for voice to the flexible, convergent, multimedia next-generation networks and the service business evolution from a monopoly-based single network to a entrepreneurial, multiservice, customer-driven environment.

In the converged world, enterprises define themselves not by their network, but by the needs of the customer. The convergence looking at all of the above is very complex and multifaceted, encompassing the domain (information technology (IT) and telecom), networks, and technologies.

Convergence is driving the radical changes in telecommunications systems and services. In the passage to converged, multimedia new-generation networks, very few legacy network systems will survive. The systems that will prevail will be those with open architectures supporting multiple technologies not only for optimization, but also for reconfiguration.

The new software systems enabling convergence and supporting multiple emerging technologies should be adaptable without being completely re-engineered. They should combine the qualities of a voice network and the opportunities of data network and provide service providers with efficient scalable growth, allowing economical expansion of the networks through the simple addition of services/up-gradations. In addition, they should enable networks and the services to be interactive and interoperable.

The intelligent network (IN) is one such interactive medium that enables users to access a large range of applications and services. It could be possible to redefine this technology to provide the consistent framework that defines the relationship between converging networks and services. This paper attempts to address one such IN technology component—the service control point (SCP). The following sections address how SCP could be redefined or re-engineered to work as a multifunction application server.

The service plane in the conventional IN could be provided as a service-oriented abstraction of network resources. The services can be considered as single distributed software objects. The Java and common object request broker architecture (CORBA) technologies would provide an advantage in building such systems.

In this suggested architecture, the service plane hosts service objects. These objects may be shared across multiple service groups. Connectivity to the network is established through the network plane. The network plane hosts network or domain-specific objects and communicates with other entities in the specific network. The architecture diagram is show in *Figure 1*.

Architecture

The architecture of a multifunction application server could be the evolution of the SCP architecture as described in International Telecommunication Union–Telecommunication Standardization Sector (ITU–T) Q12XX standards. The conceptual IN model could be re-engineered to be a more generic, network, or domain transparent model. The refined architecture comprises two major planes: the service plane and network plane. These planes interact through a set of events. These are briefly described in the following paragraphs.

FIGURE 1

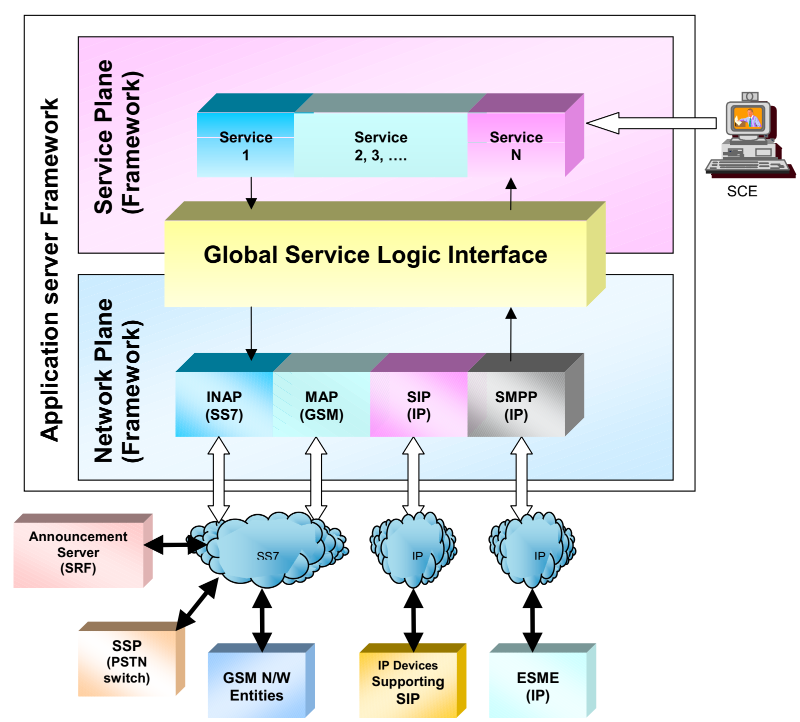

Service Plane (Framework)

The service plane and global functional plane as defined by ITU–T Q12XX could be re-engineered to a domain-independent service framework. This framework would host and execute the services transparently in the converged network environments.

The service-plane design logic is a combination of service-independent building blocks (SIBS). The ITU–T Q12X3 SIBs are re-engineered to define the services in converged networks.

The SIB is a standard reusable component having a discrete, non-interruptible operation. The SIBs can be implemented as Enterprise Java Beans (EJB). A Bean builder framework can be used to integrate and deploy the same. This integrated Java Bean logic is termed *service logic*, as shown in *Figure 2*.

The service logic may be implemented as distributed objects (using CORBA/EJB technologies) that could be shared across the service-provider boundaries, providing the distinct advantage in terms of load sharing and redundancy of services across the network. This functionality could be

FIGURE 2

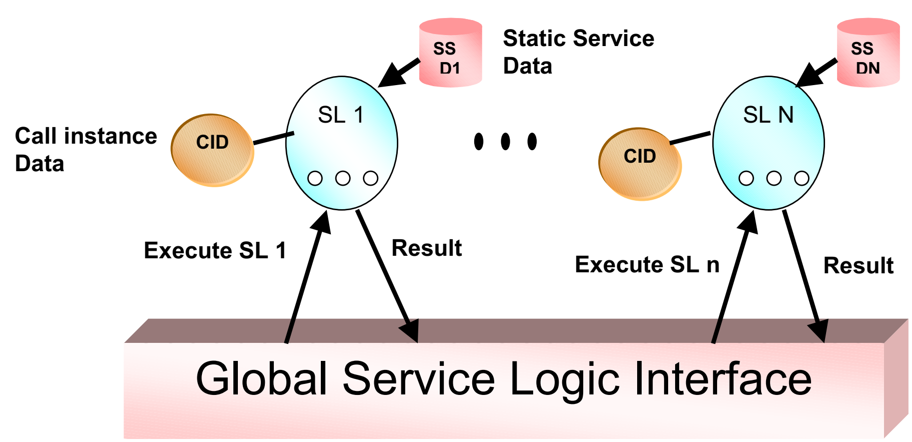

achieved through use of the service-creation environment (SCE). This design tool supports the service design across domains and protocols. The output of the SCE would be a service-logic script that can have a generic format in form of extensible markup language (XML)/Java, deployable on the application server platforms.

The execution of the service is triggered by the service request from the user on the network. The global service-logic interface enables the service instance-specific message (CID) exchange between the service logic and the network. The global service-logic interface supports the simultaneous operations of multiple services across the network domains enabling multiple services being executed from the single platform. Global service logic interacts with the network through a network plane supporting multiple protocols.

Network Plane (Framework)

The network plane is the enhanced version of the distributed and physical plane as described in Q12XX. The architecture of the framework, as shown in *Figure 3*, would support the incorporation of multiple protocols such as signalling system 7 (SS7), short message peer-to-peer (SMPP), session initiation protocol (SIP), customized applications for mobile network enhanced logic (CAMEL) application part (CAP), mobile application part (MAP) (Global System for Mobile Communications [GSM]), H.323, etc. These stacks in turn provide the connectivity with the respective network. The message converter converts the messages from the stack into the generic events understandable by the service logic and vice versa.

This type of architecture enables flexibility to add or delete any protocol interface with incremental efforts.

Advantages of Multifunction Application Server

- Services interoperable across network domains
- Easily adaptable to changing network needs
- Open framework supporting service design as per customer requirements

Applications and Services

Tele-Voting Service

The tele-voting service to support subscribers from Internet protocol (IP) and PSTN domains can be implemented on the multifunction application server. The PSTN subscribers access the service through intelligent network application part (INAP) protocol whereas subscribers on an IP network access the same through an SIP interface.

Short Message Service Center (SMSC)

The multifunction application server can be used to provide the SMSC functionality in the mobile network by interfacing with the SS7 and IP networks through the MAP and SMPP interfaces respectively. The application-management function provides the interface to database management.

Gateway Application

The multifunction application server can work as a gateway by converting the messages of different protocols. The translation of messages is implemented as service logic in the application server.

Home Location Register (HLR)

The application server can serve as a HLR toward a mobile network using SS7 protocol. HLR procedures can be implemented as service logic, and an entire HLR application would consist of coordinating service logics. The application-management function can provide the database interface.

Acknowledgements

All colleagues in Centre of Excellence–Intelligent Networks

References

1. ITU–T Q12XX standards
2. *Annual Review of Communications, Vol. 53* (2000) presented by International Engineering Consortium

FIGURE 3

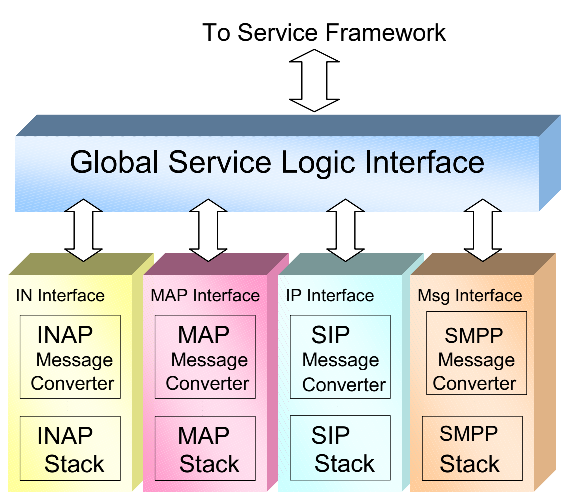

VoDSL Market Overview

Claudia Bacco
Executive Vice President
TeleChoice

This paper presents an overview of the digital subscriber line (DSL) market and focuses on current video over digital subscriber line (VoDSL) status and vendor efforts.

North American Deployment Status

The actual numbers for the second quarter show that the incumbent local-exchange carriers (ILEC) are still very much in the lead in the deployment of DSL, with residential customers being its primary users (see *Figure 1*). It should be noted that many of the service providers, if they are connecting to a business that is physically in a house such as a telecommuter or a home office, count it as a residence and not a business. These numbers are therefore not substantially skewed, but there are probably more businesses in the total number than these represent.

The first quarter of 2001 was actually the first that showed any overall slowing in the growth of DSL. One reason for that is the collapse of Northpoint. Not all of the Northpoint customers were able to reconnect with another provider, and not all of those that did were able to do it within the same quarter. So there was a downturn based on the number of people that actually went without service. The service providers that implemented price increases in the first quarter also showed a comparable percentage reduction in DSL adoption. Where prices increased there was definitely a correlation to the number of customers taking service, especially in residential service.

DSL Market Size: United States

At the end of 1998 there were 39,000 DSL lines. That number grew to 500,000 in 1999 and almost 2,500,000 last year. Almost 4,500,000 are projected by the end of this year, and 13,900,000 by the end of 2004 (see *Figure 2*).

It had been projected that DSL will surpass cable during the next year, but with this year's stumbles in the DSL marketplace, revised projections have cable remaining in the lead for the foreseeable future. The digital loop carrier (DLC) issue needs to be resolved. There is also an expectation that there will be VoDSL and value-added services (VAS) available that drive more adoption of broadband with the cable-modem data over cable service interface specifications (DOCSIS) 2.1 standards. This is going to level the playing field for cable, which might be playing more actively on the business side.

Is the DSL Market Slowing Down?

Deployments in the first quarter of 2001 lagged behind the fourth quarter of 2000. Competitive local-exchange carrier (CLEC) deployments were greatly reduced because of the Northpoint and other failures. ILEC deployments were lower at SBC and Verizon, probably linked to price increases. Interexchange carrier (IXC) deployments were still miniscule in the DSL space, compared to ILEC.AT&T has yet to see any deployment out of the Northpoint acquisition of the equipment and co-location space. And, although Sprint and Worldcom have DSL deployments, they aren't very large. Customers are feeling insecure about buying service from CLECs, and a direct linkage between the service offer and customer needs is still lacking.

Businesses require quality-of-service and performance guarantees. Before businesses jump on the DSL bandwagon, it is really necessary to make it more of a business-class service by offering different qualities, different classes, and different performance guarantees. On the residential side, service providers are starting to move beyond the early-adopter mode for DSL. Residential customers do not understand why they need DSL today; a great deal of education still needs to go on in that marketplace. Most of the service providers still talk about speed when they talk about advertising a DSL server, but content is necessary to drive the bandwidth requirement. Service is faster than dial-up, but the average residential customer does not understand why he or she needs to spend more money to go faster than dial-up.

The DSL Market Is Changing

The ILECs continue to dominate. They had 76 percent of the lines in service during the second quarter of last year and 84 percent this year (Obviously, Northpoint had a big effect on that number). ILECs grew 16 percent in the second quarter, while CLECs grew only 6 percent. Again, Northpoint drove most of that change. Regional ILECs are still doing quite well with advanced services. But the CLEC shakeout is, unfortunately, probably not over. Although Northpoint, Rhythms, and many others are gone, Covad is struggling to exit from its Chapter 11 filing. IXCs are starting to be more aggressive, but there is not a huge deployment in place yet. AT&T still has not acted on the Northpoint acquisition, and Worldcom has just announced the purchase of key Rhythms assets and has some limited DSL assets from UUNET. Other

FIGURE 1

North American Deployment Status

TeleChoice 2Q01 DSL Deployment Summary			
Service Provider	2Q01 Lines in Service	% Residential	% Business
ILECs – US	2,800,160	80%	20%
CLECs – US	508,331	42%	58%
IXCs – US	26,000	12%	88%
Total	***3,334,491***	***74%***	***26%***
Canada	***780,943***	***76%***	***24%***

CLECs and integrated communications providers (ICP)—NewSouth and New Edge, for example—have actually backed away from DSL and are focusing more on virtual private network (VPN) services. Smaller players, like data competitive local-exchange carriers (D–CLEC), are leaving the DSL market.

VoDSL Status

What does all that mean for VoDSL? Although there are deployments, there are still not any major ones, and everyone is waiting for the coming AT&T deployment. But there are many small deployments today; a lot of regional players are doing well offering VAS and VoDSL services. Regional CLECs are doing quite well by offering very tailored packages to the small and medium-sized business market, but they are in very small pockets of the country.

Many people overlook the regional ILECs, the very small incumbent players, but they are really the most progressive in the marketplace. They have VoDSL deployment, and there is a surprising amount of VoDSL deployment for entertainment by small regional ILECs that are serving the customers with very sophisticated services. Granted, though, these are players whose entire market might be 2,000 customers and 2,000 lines.

Among the major players in VoDSL deployment over the last year, Picus, a CLEC provider, has gone out of business; mPower—the leader today as far as VoDSL currently in service—has announced a reduction in the number of central offices (CO) that it is serving; and Trivergent merged into NuVox and is only offering voice over T1 (VoT1).

VoDSL—Why So Slow?

Why is everyone moving so slowly? Small players are doing small deployments. The D–CLECs have not done anything. One thing holding up deployment at players like Covad, Northpoint, and Rhythms was not the technology or the functionality, but the fact that they had to partner with someone else for a voice network and they were struggling with who charged what, how to divvy it up, and how to bill the customer.

The ILECs are all doing lab trials, but none of them have been really aggressive in moving toward deployment. ILECs are very interested in VoDSL for their out-of-region CLEC operations and in areas in region where they have copper exhaustion, but they still are not really thinking about VoDSL as a normal part of their deployment. However, cable telephony is going to become a much more viable competitor for them, and they are realizing that there are 1,000,000 cable-telephony lines in service in the United States. That number is not huge compared to plain old telephone service (POTS), but for voice over cable that is a pretty impressive number, given the number of cable lines in service. They are definitely making a lot of headway there. There are primary-line services and secondary teen-line or business-line services in the mix. People are not only doing a second-line scenario; there are people adopting a primary-line scenario over cable.

Vendor Efforts

What are the vendors doing to move DSL along? For one, they have improved gateways to the point that the ILECs believe that they are truly "carrier class." They are meeting the network-equipment building standards (NEBS) compliance for scalability and redundancy at the level where a large service provider feels comfortable. There have also been a lot of standards bodies put in place, including the open voice-over-broadband (VoB) effort and interoperability labs. Customer-premises equipment (CPE) has improved, and the prices of integrated access devices (IAD) have come down. They are still not low enough for residential use, but they are certainly moving in the downward direction. The Panasonic/Jet Stream announcement that Panasonic is getting into the VoDSL business makes a really strong statement about acceptability in the consumer marketplace. Panasonic is a very strong consumer-oriented brand that people feel comfortable purchasing, and for them to have a VoDSL IAD product is a statement that the technology is ready for primetime in the consumer marketplace.

Panasonic's VoDSL IAD is very consumer friendly. It looks and works just like a normal phone—the bottom is just a little thicker. People will not be put off by having it in their house. The gateway providers are also hedging their bets—no one is really talking about VoDSL anymore. Instead, there is VoB, and there is also a lot of conversation about moving to voice over Internet protocol (VoIP) and partnering with softswitch vendors so that all the bases are covered.

Who Is Going to Make the First Move?

Does VoDSL come out first? Or does VoIP? Or do they come out together? What is coming next?

Most of the potential VoDSL providers are also talking about providing the IAD as part of service, as opposed to selling it today, because of the price point. But to get to a retail model, where customers could buy an IAD at a department store, the price needs to come more in line with the standard subset. Indeed, it could be argued that Sprint ION is VoDSL. It is VoIP over DSL. In fact, if there is not a gateway involved, that is really VoIP technology. It is a proprietary CPE technology, but the connection to the home or the business is over DSL for the consumer market. On the business side, there can be different transport technologies, but on the consumer side it is all DSL.

The ILECs are still moving slowing, though some vendor selections are expected by the end of the year. And even after vendor announcements, long technical and market trials are anticipated.

Or Maybe Nobody?

According to most ILECs, the quality is not there for VoIP to be a reality in the near future; it will probably be three to four years before it is a real service offering with the necessary quality of service. Between now and then, there would just be very small pockets of early adopters that would do it, or businesses within their business on a campus.

If VoDSL does not happen soon and VoIP continues to progress, will we have a VoDSL? VoDSL has a limited "shelf life." If deployments do not come before VoIP solutions are equally mature, they may not come at all.

There probably will be a migration. True VoDSL, where there is a gateway and DSL is the transport, would be the

FIGURE 2

DSL Market Size: United States

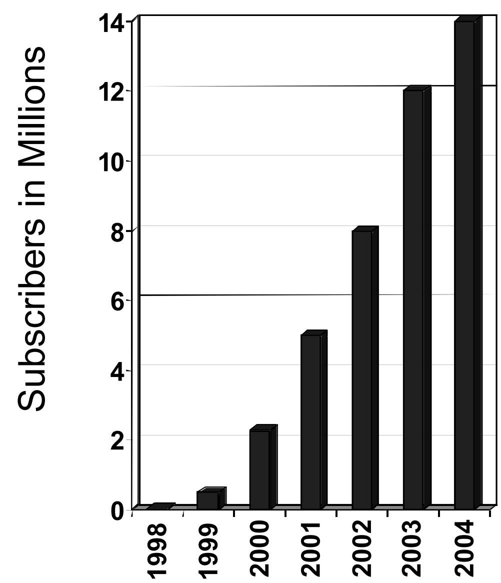

Source: TeleChoice, Inc. (January, 2001)

first deployment. Then there would be many levels of VoIP, where there would originally be VoIP over DSL, but it still goes through a gateway and actually travels over Class 4 and 5 networks. An entire VoIP infrastructure is really many, many years away. Service providers will start with VoDSL and progress over the long term before truly getting to VoIP. Although some people say that if VoDSL does not happen soon, it may not happen and we will jump right to VoIP, it is not likely that VoIP is going to get through its issues any more quickly than VoDSL.

Of course, this discussion is specific to the United States. Outside the country, there is much more willingness to adopt VoIP–VoDSL type of services. They will probably take off and be adopted much more quickly outside of the United States.

Service providers, for the most part, are still very limited in their focus. ILECs are more likely to pursue gaming, security, video, and other "entertainment" applications first to counter cable companies. For example, Speakeasy.net, an Internet service provider (ISP) in the Seattle area, is residentially focused, but it has packages for specific target markets. It has the gamer package, which has guarantees on latency and number of hops; it has a day-trader package; and it has a system administrator packet—all this and with almost no advertising. It is all word of mouth within these communities, but it has become incredibly well known for having the services to fit these very tacky little cliques within the industry. It is a provider that, from a VAS perspective, can be held up as an example in the residential market.

Small businesses are still underserved, though there are regional CLECs that are really striking out at the small-business market. ILECs are still, for the most part, ignoring the small-business market. They are just starting to realize that this can be a big market for them. There is a lot of opportunity in the small-business market for service providers thinking of going with VoDSL, and there is a huge opportunity in the small and medium-sized enterprise (SME) marketplace. They are currently pretty much ignored by the large providers. They have the low-cost residential services that probably do not have the functionality they need or the larger business services with functionality but with a price point that is way out of their range. So there is not a lot out there today that is targeted for that market.

What Do DSL Providers Need to Do Next?

Right now DSL service providers are focusing on profitability—everyone said that this year was going to be the year of VoDSL. Why didn't it happen? Is there something wrong with VoDSL? No, because last year the service providers were focusing on deployment at all costs. Then, at the end of the year, the market suddenly wanted to know where their profitability was. So this year all of them are focusing on being profitable on the core DSL services and VoDSL. VASs have not been taken off the table; they have just been delayed while providers focused on profitability.

Therefore, profitability is clearly number one. Number two is for the ILECs to really recognize the competition. Among ILECs there has never been a real fear of the CLECs taking away any substantial part of their business, and with a lot of the CLECs crumbling there is even less fear. Cable is serious about voice and is turning into the real threat for the ILEC market. The IXCs are starting to look like the more progressive players for the large players. There is an opportunity for the IXCs to make a serious attempt at bringing voice and data into ILEC territories and really learning from the small guys. There are a lot of smaller players like Speakeasy.net and Telocity that are very successful with targeted packages, creative, value-added packages for smaller audiences. Thus there is an opportunity for the large players to think like small players and really get ahead in the marketplace.

Security Risk Management: Challenges for Today and Tomorrow

Charlie Baggett
Chief Technology Officer
SAFEOperations, Inc.

Abstract

This paper provides a quick overview of the state of information security, taking a fairly critical stance at times. It is helpful to remember, however, that this is a very young field, and most of the disciplines associated with it are immature. This paper focuses on the motivations of information security, reviews the state of the practice, evaluates emerging standards, and offers some solutions. Finally, it considers what the future will look like, with particular emphasis on the necessity for moving beyond signatures to predictive analysis.

Motivation

The biggest cause of concern about information security is reaction to a serious incident. That is not surprising because once users are harmed in some way, they turn their attention to security and become much more interested in the larger issues. What usually follows is concern about the consequences of rapidly expanding connectivity, especially mobile computing (landline and wireless), that, while certainly aiding the growth of business, introduces a new set of problems in itself. High-speed digital subscriber line (DSL) technology in homes (and hotels) facilitates telecommuting, while the wireless world deals with personal digital assistants (PDA), laptops, and cellular phones. Another added layer of complexity is virtual private networks (VPN). Thus, all the problems introduced by those technologies are introduced into the entire network. They lack identification and authentication, with no intrinsic encryption, so users must employ add-ons.

Customer service is another motivating factor behind information security. Businesses want to provide more information and gain more access to their customers, but at the same time they want that information to go only to the customers. Businesses do not want that information going to other people, particularly competitors. Other underlying motivating factors include network and software complexity, the human element, security as "bolt-on" in business and application planning without a systems worldview, and the shortage of security professionals.

State of the Practice

Business Recognition
Businesses must first recognize that senior management must pay sufficient attention to security problems, particularly in their information systems, and then provide the resources to solve those problems. From this recognition flow the policies and procedures to ensure the information security implementation that tracks the goals and objectives of the business. To accomplish all this requires careful attention to staffing and training. However, awareness training is needed for the entire workforce—not for only the security practitioners—because there are a number of weaknesses in any organization.

Technology
The biggest set of technology issues is associated with configurations. Most of the malicious activity today, particularly the scripted schemes, is directed at configuration problems and known software bugs that are not patched. Scaling issues could fill an entire briefing, but one particular issue deserves attention. Private public key infrastructures (PKI) are desired by businesses that want to have public keys that are not accessible to everybody. In other words, they want some privacy, which introduces a registry security issue into the whole set of PKI issues.

One problem that pervades information security involves user authentication—being able to identify who is actually interacting with the system. Current practice is for users to choose passwords and passphrases, but this is a very weak approach to system security. The short development life cycle of security products contributes to system insecurity because enhanced feature sets are rapidly added to already immature security products before existing problems are identified and solved. Therefore, consumers are the problem finders. In other words, the users find the bugs and then report them back, with the expectation that they will be fixed. The industry needs to figure out how to get beyond that through improved operating systems (OS), devices, and applications.

In the misuse and incident detection area, the current products on the market are relatively crude. Most are signature based, and it is simple to demonstrate how easy it is to bypass signatures. Encryption, when used by attackers, can blind many of these products. The telecommunications industry has difficulty defining what "normal" is in networks. It does not know how to do it, but it is learning how. Yet defining "normal" involves a temporal problem because the criteria for "normal" change over time. Products are only beginning to integrate monitoring and management capability. As this trend accelerates, the outputs of those products will help security specialists to gain a better understanding of the problem. Currently, it is difficult to monitor high-speed, large-scale networks; to solve problems when there are attacks against those networks; and to discern what is going on, where attacks are coming from, and how they are composed. Network security today is in a reactive mode. Things happen after the fact, which is very costly because it requires paying for investigations. In addition, the loss of business and reputation associated with any attack can have devastating economic repercussions. Worse, because current security paradigms deal with attacks after they have been identified, the full extent—and success—of an attack might never be known.

Emerging Standards

One of the key emerging standards will be the Internet protocol (IP) version 6. Designed with security in mind, it has a 128-bit address space, authentication, encryption, host interface address switching, and numerous other new features. In the meantime, though, with IP version 4, IP security (IPSec) has been implemented, and that requires the implementation and deployment of some kind of PKI to tighten user identification and authentication. Another area of weakness is in the domain name server (DNS) area, where a requester makes a request and has to believe the response. Unfortunately, it is easy for someone to spoof the rudimentary authentication features. For example, caches are vulnerable to unauthorized access and insertion of incorrect or malicious data. Undetected software deficiencies can subject the network to simple attacks with disastrous consequences. Servers are also subject to attacks, such as denial of service (DoS) and buffer overflow attacks.

Coordination among providers is a growing concern. Most of the attacks that occur now take place across providers and across legal jurisdictions, requiring multiple people to solve the problem. Even if providers agree that they must share information, they still have numerous questions, such as: What do they share? What format do they put it in? How do they find information that needs to be shared in this huge sea of packets? It is not a simple problem. Because attacks cross jurisdictions, law enforcement organizations face a plethora of legal, political, and technical problems. When attacks originate internationally, these issues are magnified, causing problems with attaining and executing court orders. But law enforcement has been pretty slow getting up to speed, both in the technology and in the legal process aspects of this. Law enforcement is complicated by the industry's reluctance to share and expose information on attacks because providers are concerned about what will happen to that information.

Solutions

There are several steps to take to secure an environment, although there are basically no current system-wide solutions available. The first step begins with a risk assessment, a very expensive task performed largely by consultants. The results can be inconsistent, not because of the quality of the work, but because there is no solid baseline against which the assessor can measure. It is hard to make measurements over time and from peer to peer. The next step involves policies and procedures, the road map to information security. Because of the lack of attention by senior management, consultants often find that policies are either incomplete or nonexistent and that the ones in place are usually not enforced. Worst of all, because no policies are in place to track the business, business policies do not drive the selection of technology or the implementation of security solutions.

An in-depth defense is needed. Attacks have occurred at all layers of the protocol stack. At the network level, guarding the information assets includes the use of routers, firewalls, and antivirus software. One of the first things needed in any system is strong authentication that is stronger and easier to implement than the current password system but not as costly and expensive as implementing a comprehensive PKI. Something is certainly needed now to fill this gap.

Products must be centralized so that they can be centrally monitored, which requires more vigorous product development and testing. More proactive monitoring is essential. Information should be consolidated and correlated across various aspects of the OSs, not done piecemeal. Continuous feedback is important, and the monitoring and analysis systems must be updated. A vigilant monitoring system will have a network intrusion detection system (NIDS), a host intrusion detection system (HIDS), and limited application misuse detection. The industry must move away from strict signature identification and should learn how to monitor in high-speed, large-scale environments in real time. Instrumentation of the entire computing environment is needed to attain more complete network security.

Network security involves a very complex set of issues. To tackle the problem, service providers may choose to outsource information security for their networks. The managed service provider should not be somebody who does it as a sideline; instead, experienced security professionals should be hired who have a demonstrated history of providing thorough, timely security services. In addition to the experience of its employees, the service provider should have a robust, secure operations center from which to manage client networks.

Beyond Signatures: Predictive Analysis

Any network-security solution must work at the network, host, and application levels. A coherent solution will address the entire stack simultaneously and will allow the establishment of norms in preventive action so that damage to the network can be limited. Predictive analysis will allow network security to move beyond signatures in the current intrusion detection systems. Processing must go as far forward in the system as possible. Effective instrumentation must be included at all levels of the network. Products are now opening themselves to the capability for monitoring and management. Ideally, this would be done as transparently as possible, without having to place sensors in the networks, and using the minimum necessary bandwidth. The establishment of network-specific behaviors, norms, and patterns will help predict attacks and malicious behavior. Those behaviors, norms, and patterns specific to the traffic flowing over the network can be linked to the geographic locations in which activities are taking place and to the people who participate in the conduct of the business so that security analysts can begin to understand these parameters as a function of time. The data must be timely and reliable. The false-alarm rate must be compatible with the business objectives, and the tools must be able to learn and to keep up with Internet time. Predictive analysis will support the focusing of resources, which will have benefits beyond merely the capabilities to handle security items.

Leveraging Optical Options

Richard Barcus
President and Chief Operating Officer
Tellium

Incredible advances in technology have occurred over the past couple of years. This is especially true in the world of communications, where advances in optical networking have made new and exciting broadband applications possible. While the emerging opportunities show promise for the end user, carriers are struggling to support these applications while remaining profitable. This paper discusses leveraging the optical layer to help carriers introduce new broadband services while remaining profitable.

Carriers were once enthusiastic about reports of bandwidth demand doubling every six months. Incumbent carriers raced to expand their networks, and new carriers built an infrastructure from the ground up. Fierce competition drove the expansion of infrastructure, but it also drove bandwidth prices into the ground. Carriers began to question their ability to respond to the demand while also making a profit. As bandwidth prices plummeted and marginally profitable carriers began closing their doors, a number of reporters and analysts began to cite a bandwidth glut. The idea of a bandwidth glut appears ludicrous, though, given that most carriers continue to struggle to meet demand.

Nevertheless, these carriers do recognize that building and maintaining an infrastructure is expensive. Their current drive to build a robust infrastructure that can handle the emerging applications has their costs rising at a rate that outpaces the corresponding increase in revenue. Consequently, these carriers are now looking for a way to decrease costs, increase revenue, and create a sustainable competitive advantage over the competition.

While the short-term future of the industry remains a challenge, the long-term possibilities remain extremely encouraging. Success will be achieved by those carriers that recognize that the communications landscape has changed. These carriers must plan for the long term; they can no longer forecast and build infrastructure based on the traditional traffic-engineering model using the predictable and steady growth of voice traffic. A number of emerging technologies support a multitude of broadband applications, which could explode and turn a perceived bandwidth glut into a real bandwidth shortage overnight. Remote-office networks, WebMD, instant messaging, Internet gaming, distance learning, webcasts, MP3 audio, streaming video, electronic photo albums, e-libraries, e-movies, and on-line museums are just a few applications currently driving the need for bandwidth. The challenge for carriers is to make these applications available to global consumers at a fair and affordable price while remaining profitable. On-demand access with a pay per use creates a need for a dynamic network. Bottlenecks, which could kill an application, must be eliminated in both the core and edge of the network. Eliminating bottlenecks in the core of the network has been the focus of long-haul carriers for the past couple of years and includes a fairly well-defined solution.

Eliminating the bottleneck in the metro—particularly the first mile where end users request large files or listen to music over the Internet—remains a challenge. A host of companies are using a wide variety of technologies to address the problem. These include cable modems, satellite, private networks, fiber to the home, synchronous optical networks (SONET), new-generation SONET, asynchronous transfer mode (ATM), Internet protocol (IP), metro dense wave division multiplexing (DWDM), or the all-in-one multiservice protocol platform. The lack of a clear, winning solution, though, exasperates the problem. A clear winner would benefit the end user by allowing investment in a technology with longevity. This should, in-turn, support continued bandwidth growth in both the edge and core networks.

Technology advancements in the core provide carriers with an opportunity to create a network that is more scalable, flexible and dynamic than today's core network. These advancements are based on optical switching technology, which has had an accelerated time of development over the last couple of years and promises further developments over the next decade.

Market Influence

Advancements in technology were aided by the economic well-being of the United States during the late 1990s. The abundance of investment capital funded sophisticated and innovative developments in a number of areas, including optical switching. Small, eager firms with focused research and development organizations were able to commercialize optical technologies that had been in the laboratories for a number of years. Investment funds also flowed freely into communication carriers with aggressive plans to profit from the unprecedented demand for bandwidth. Incumbent carriers saw their territories challenged by new, aggressive upstarts as well as by non-traditional players such as cable companies and utility consortiums. End users finally began to see the competition and, more importantly, the lower

FIGURE 1

Carriers' Dilemma

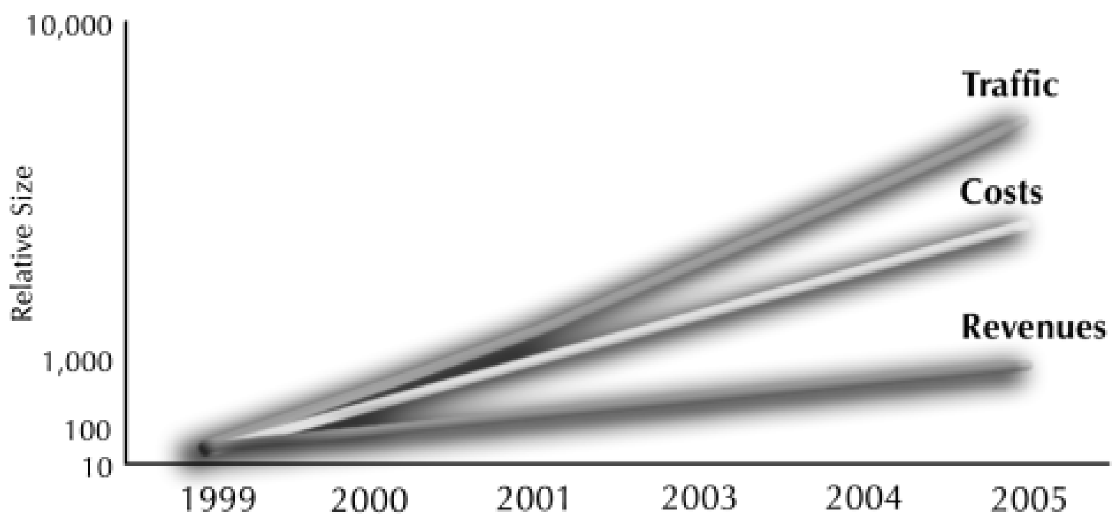

prices that were alluded to at the time of deregulation. Carriers began forward pricing their services, i.e., cutting prices to attract new customers. The result was more carriers competing for the same pool of customers and cutting their prices to gain a competitive advantage. Carriers had to build out their networks to support the anticipated demands, while dropping prices to attract new customers. This contributed to costs rising at a faster rate than revenues, as indicated in *Figure 1*. Given these circumstances, it is not surprising that some carriers had to shut their doors.

The recent turn of the market is unfortunate, but it should help stabilize pricing by reducing the number of competitors and by concentrating bandwidth demand among a manageable number of suppliers. The key to carrier survival will be to reduce the cost of building out their infrastructure while continuing their evolution toward a network that is more scaleable, flexible and dynamic. These carriers should then be able to offer new optical services, supporting the emerging bandwidth-intensive applications. Their newer, lower cost basis will permit them to price their services aggressively, provide better service, and enjoy higher margins.

Optical Network Elements

Impressive developments in key optical-networking components have occurred over the last couple of years. It is imperative that the industry builds on this momentum and maximizes the benefits of the resulting optical network elements. Continued advancements on these products are the key to carriers supporting the bandwidth growth that will, in return, encourage new applications. These network elements, as depicted in *Figure 2*, include DWDM, optical add/drop multiplexers (OADM), and intelligent optical switches. Advancements in these areas hold the promise of supporting new bandwidth-intensive applications in medicine, education, law enforcement, banking, and entertainment. Taking a quick look at each of these elements will show how continued development in each supports the evolution of the optical network.

DWDM is probably the key element that defines optical networking today. This device, which quickly established itself in core long-haul networks, solved the bottleneck issue for inter-exchange carriers who were short on fiber and facing challenging right-of-way issues. DWDM created the opportunity to increase the capacity of fiber 100 fold. Continued development of this element promises to add unlimited capacity on a given physical fiber path. DWDM manufacturers are promising to open up new transmission bands, while decreasing channel spacing to offer virtually unlimited capacity on a single fiber. While the prospect of thousands of wavelengths on a single fiber is exciting, it is the process used by DWDM that makes new services possible. DWDM transmits hundreds of discernable light signals as individual light paths. Each of these light signals can carry different bit rates, different technologies, and can be routed to various parts of the network. In essence, it provides the foundation for optical networking. DWDM is well entrenched in the core of the network, and continues to show promise in the metro to allow the extension of end-to-end optical services closer to the user.

As the number of wavelengths increases, the ability to access a subset of them becomes increasingly important. It would be overly expensive to demultiplex all of the wavelengths just to access a small subset, and OADMs have been introduced to address this issue. OADMs will play a major role in the continued evolution of the optical network. These devices are not very mature, however, and are just now emerging as reliable elements that allow access to a subset of wavelengths without converting the optical signal to electrical and back to optical again. OADMs generally use filters, which introduce loss into the network and until recently only allowed access to a very small portion of the

FIGURE 2

Optical Network Elements

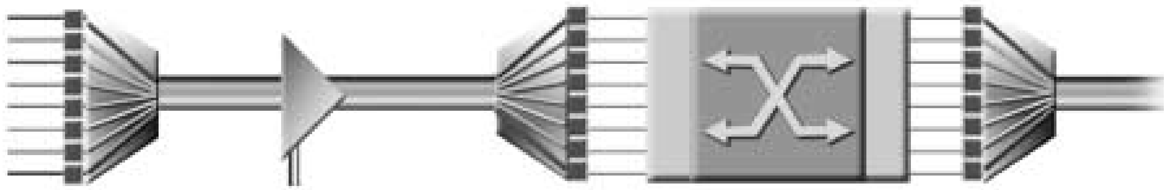

wavelengths. There are promising developments that allow access to a growing number of these wavelengths without introducing an unmanageable amount of loss. In addition, developers are working on complementary products to allow carriers to dynamically select the wavelengths they want to add or drop. The continued development and deployment of OADMs will help reduce the cost of optical networks while helping to improve its flexibility.

While DWDM and OADMs create a foundation for optical networking, the true promise of a network that is scalable, flexible, dynamic, and economically viable will not be fulfilled without optical switches. Unlike OADMS, optical switches can manage a large number of wavelengths entering a network node from multiple directions.

Optical switches provide large-scale photonic management of signals between DWDM terminals. These intelligent devices add manageability to the large number of optical signals supported by today's DWDM devices. Optical switches that normally interface with numerous DWDM terminals generally process signals at the transport-network line rate of 2.5 gigabits per second (Gbps) or 10 Gbps. These optical switches empower carriers with the ability to remotely provision these signals, rapidly reroute these signals around failures, or implement new and varied optical services. Because the majority of this bandwidth is carrying an increasing amount of data traffic, it is important to note that core optical switches, grooming at 2.5 Gbps, can perform these functions without adding appreciable network latency.

Intelligent optical switches also allow carriers to use more of their network to generate revenue. Most carriers today use dedicated ring-network architectures that reserve half of their bandwidth for protection. This results in circuits that sit idle most of the time waiting for something to go wrong. Intelligent optical switches support multiple architectures, including mesh, which will allow a greater portion of the network to be used to carry revenue-based traffic. Mesh networks, as shown in *Figure 3*, not only allow working circuits to share restoration routes but also provide optional paths for the unlikely event that the primary restoration route is occupied. The ability to dynamically select various restoration routes makes mesh networks more reliable than rings on a network-wide basis. Optical switches, deployed in mesh architectures, provide carriers with the ability to offer various grades of service. They can continue to offer dedicated restoration paths to those discerning customers who insist on, and are willing to pay for, premium services. Carriers may also offer discounted services to a consortium of users that establishes a dedicated working path but shares restoration routes. Shared restoration can be used for diverse working traffic where it is unlikely that simultaneous failures would occur. This method increases the efficiency of a carrier's network and uses more of it to generate revenue.

Flexible architectures such as mesh, which use intelligent optical switches, ultimately give the end user more control. Optical switch-powered networks also give customers control of their own restoration scenarios, invoking various schemes based on the time and type of traffic they are supporting at any given instance. This type of control can be supported through either the introduction of improved hardware or sophisticated software support. The use of software currently has the most momentum.

Sophisticated software algorithms work together as the central nervous system for the next-generation of optical mesh networks. Improved developments in this area will allow the optical switch to recognize service conditions and match them with service obligations, performance criteria, and maintenance events. The correlation and subsequent reactions will occur for either planned or unplanned events. The software will process all of the information and then take action based on criteria established by either the carrier or customer.

The Case for the All-Optical Switch

The current state of network evolution includes the deployment of intelligent optical switches, which uses an electrical core. These switches are known as optical-electrical-optical

FIGURE 3

Mesh Network

(O–E–O) because they convert an optical signal to electrical and the back to optical. While the conversion supports a number of useful functions, such as grooming and 3R regeneration, some see it as it as introducing unnecessary expense into the network. Proponents of an all-optical switch conclude that the elimination of electrical components will result in a switch that is significantly less expensive than the current generation of O–E–O. While this may be true in some limited applications, the lack of any significant deployment of all-optical technology indicates that the all-optical switch is not yet ready for prime time. Without the electrical components, all-optical switches have no way of passing important performance data. They lack the ability to communicate vital network information used for restoration and dynamic provisioning. So while eliminating the costs of electrical components, all-optical switches also eliminate the intelligence required for a more flexible and dynamic network.

Certainly, the promise of the all-optical switch is well understood. Theoretically it is infinitely scaleable, supporting any bit rate on each of its ports. Thus, the throughput of the switch can be greatly increased simply by increasing the line rate of the subtending transport equipment. A 256 port all-optical switch can support growth from less than 1 to 10 terabits throughput by upgrading the transport equipment from optical carrier (OC)–48 (2.5 Gbps) to OC–768 (40 Gbps). If the switch is truly transparent, this is accomplished with no impact. However, as the number of ports increases, so does the loss through the fabric, the complexity of control, and the difficulty of fiber management.

All-optical networks require a great deal of sophisticated optical engineering to plan for physical impairments that accumulate over an end-to-end path. Such preplanning is, of course, counter-intuitive to the movement toward a more flexible and dynamic network. The introduction of an all-optical switch into an all-optical network may actually cause the total network cost to increase by requiring additional engineering and equipment to accommodate the different optical performances introduced by the appearance of a new optical element to the network. An all-optical network currently has to be engineered for worst-case scenarios, thereby eliminating the savings that one might otherwise expect from optimized network links. In addition, all-optical networking requires the implementation of equipment from a single vendor and could eliminate savings garnered from using best-of-breed for particular areas of the network. Open standards, which support interoperability at the optical level, do not exist and are not expected in the foreseeable future.

All-optical switches, implemented correctly and at the right time, will further the cause of optical networking. The economics of an all-optical switch in a core network doesn't appear to be viable until an appreciable number of high bit-rate signals (40 Gbps and above) are deployed. In addition, these signals will support optical services that do not require grooming at intermediate nodes. Although it looks like wide-scale use of 40 Gbps transmission will be a market driver for all-optical switching, there will be a real need for grooming 2.5-Gbps and 10-Gbps payloads within these large 40-Gbps signals.

Given the state of all-optical switches and carriers' continued need to groom information from data routers operating at 2.5 Gbps and 10 Gbps, O–E–O switches will be the primary mechanism for managing wavelengths for many years to come, and all-optical switches will be introduced into the network to complement the functions of an O–E–O switch. All-optical switches will support high-speed bypass at a network node, while the O–E–O switch will continue to support functions such as performance monitoring, fault isolation, grooming, and restoration. Together, rather than separately, they represent the suite of advantages needed by carriers to build networks that can provide the services, security, and reliability that will lead to profitability. O–E–O and all-optical switches are not competing; rather, their strengths complement one another. As complementary technologies, they provide carriers with a painless way of using either or both based on the application.

Looking Ahead

Carriers can look forward to continued advancements in optical switching technologies. These advancements, coupled with a sound business plan, will return profitability to the communications industry. The deployment of intelligent optical switches, coupling the latest in hardware development with sophisticated software, will allow carriers to offer new optical services while decreasing provisioning times and lowering both capital and operating costs. Optical switches offer a realistic approach to evolving the network to support new services, without disrupting sound operating processes and procedures. Intelligent optical switches help carriers recover valuable floor space, lower power consumption, introduce new services, and offer differentiated restoration services. They take advantage of advancements in electronic components to support continued scaling, while supporting the maintenance functions required of core optical networking. This evolution will consist of the addition of all-optical switches, supporting the dream of bit-rate independence while offering unlimited scalability. All-optical switching will be added to network nodes and work in conjunction with O–E–O switches to maximize the benefits of both bandwidth and optical services; meanwhile, the O–E–O switch will continue the vital role of grooming and providing valuable performance monitoring information. The all-optical switch will efficiently process the bypass traffic, which neither terminates nor requires grooming. Together these switches will provide increased flexibility without sacrificing the carriers' ability to localize a particular fault down to the card level.

Significant progress is expected over the next couple of years to extend the best-of-breed optical switching concept to include the use of a common control mechanism between not only optical switches but also data routers. This will provide the impetus for new applications and a further reduction in provisioning times. In the end, the clients in many cases will be able to turn up, turn down, change, and alter their own capacity when and where they desire. This function will be supported by network management systems that interact with the optical switches to handle billing functions.

Most of these functions, driven by demand and factors of economy, will occur in the core of the network first. However, as technology matures and the direction of the edge market becomes clearer, optical networking will move closer to the end user. Further development at the edge and core, particularly in the area of control, will make true end-to-end opportunities a reality.

Where Communications and Computing Meet

Craig R. Barrett
President and Chief Executive Officer
Intel Corporation

We're going to talk about the digital world. We'll talk about the divergence of communications and computing and try to give you some historical insight on what's gone on in the past, what sort of evolutionary changes there are, and what we need to do.

So we'll talk a little bit about the state of the industry today. We'll talk about the principles of technology evolution, about how you get from one generation to the next, and a little bit about our industries. I want to emphasize "our industries." It's really the computing and communications industries working together to bring solutions to the consumer, whether the consumer is the home consumer or the business, a large business, or a small or medium-sized enterprise.

We've all been impacted by an economic slowdown; that's no secret. There's a dot-com meltdown. There's a carryover from Y2K. Information technology (IT) spending has slowed down dramatically, especially here in the United States. Yet it's still going gangbusters in other geographies around the world. But purchasing of our equipment, some of your equipment, has slowed down. The mobile handset market has slowed down dramatically. It's still growing, but with nowhere near the growth that it had a year or so ago.

If you look more particularly at the communications industry, it is an industry that invests a lot of money. Debts are at record level. You find people who are bidding billions upon billions of dollars on third-generation (3G) spectrum for the next-generation handset wireless capability.

The estimate is that it will take something like $150 billion to buy the spectrum in the various auctions run around the world if those auctions continue to take place. One was canceled in Singapore recently for lack of interest. But $150 billion to buy the spectrum, $150 billion to put the infrastructure in place to deliver that capability, $150 billion to provide the applications to get the consumers to come before you start to see a nickel of return on that investment.

So there's a huge challenge in some of the communications areas for the next-generation capability.

There's also some degree in excess of capacity with dark fiber underground. We've looked at what's happened recently to some of the networking communications suppliers, the original equipment manufacturers (OEMs) in this space. We've seen that it grew like there was no tomorrow and that the business would never slow down. It would grow on at 50 percent compound annual rate going forward.

Well, it turns out that supply-line management wasn't wonderful in that space. As orders fell off, we've seen massive inventory write-offs. The entire supply chain has been impacted to various degrees.

That is the state of the industry today, but it's going to get better tomorrow. If not tomorrow, then next day. And the real reason is that we have a digital world ahead of us. And there's no doubt that this digital world is not going to slow down. If you look at the number of Internet users, the number of computers, the number of personal digital assistants (PDAs), the number of handsets that were sold each year, those numbers continue to go up.

We know that the Internet buildout is going to happen. We know that the Internet is going to be the medium for communication, for entertainment, for access to information, and for conducting business and commerce. We know that that's going to continue.

Every two minutes, there are approximately 50 new subscribers to the i-Mode phone at DoCoMo, which is the always-on packetized capability with handsets in Japan.

There are approximately 400 new users joining the Internet around the world every two minutes.

There are approximately 1,400 new auctions listed on eBay every two minutes. And we're just going to continue scaling this up.

About 1,500 cell phones are sold every two minutes. That translates to roughly 400 to 450 million new cell phones a year.

Amazon.com sells about $11,000 of goods over the Internet every two minutes.

And look at how many searches are completed at Google's high-performance search engine: something like 83,000 searches every two minutes.

And in its heyday, there were something like 100,000 file transfers started every two minutes on Napster.

If you look at how many instant messages are exchanged on America Online (AOL) every two minutes, it's something like 900,000 short messages.

Somewhere between 500,000 pages and one million pages are downloaded from Yahoo! every two minutes. This continues to grow.

And there are about 50 trillion voice bits on the backbone every two minutes. And if you look at data bits on the backbone, there' something on the order of 100 trillion data bits every two minutes.

Every one of these numbers, save Napster, perhaps—but there will be follow-ons to Napster—will continue to increase as we go forward. This is really the future of our two industries: This immense amount of information being communicated over the various backbones and then the computing power to be able to analyze it, manipulate it, deal with it, redirect it, enhance it, and enjoy it.

One of the things that bother me a lot is reading in the press about the United States and how our Internet buildout is complete, mature. The United States has about 5 percent of the world's population. And if you look around the world, you're starting to see immense activity going forward in terms of number of Internet uses, users, and capabilities.

Europe, for example, is projected to have more Internet users than the United States within the next one to two years. They have the highest mobile phone usage in the world: more than 52 percent penetration of mobile phones—higher than that of Japan, higher than that of the United States.

In Asia, Internet usage is growing by a compound annual rate of about 45 percent. They're projected to have 400 million users on the Internet by 2004. Within a couple of years, Asia (excluding Japan) is projected to have 25 percent of the global Internet users.

Meanwhile, although Japan has discovered the latest in cellular phones, packet switch capability, and the i-Mode phone with 20 million users, at the same time they've also grown their Internet usage phenomenally: more than 60 percent compound annual growth rate (CAGR) in the last year or two as the Japanese consumer has discovered the Internet. Furthermore, about 50 percent of the Japanese consumers are, in fact, mobile phone users as well.

The key issue to all of these statistics is that the Internet is growing dramatically and will continue to grow dramatically—and that's what we're all about: using the Internet for communication, information access, entertainment, and conducting commerce. And that's why we collectively need to work together—to make sure that that Internet usage continues to grow and that the world becomes increasingly digital.

So the Internet is our growth engine. And whether we like it or not, we're both—the computing industry and communications industry—tied, together, to that growth engine.

And if that is the growth engine, then it probably makes sense to look a little bit at the principles of that growth. What do we know from the past about how technology moves forward? What can we do about it in the future? What can we expect in the future?

We're really starting off today in the communications space with three separate networks. They're relatively inefficient today as three separate networks, but as we move forward, the inefficiencies have to grow out of the system, and the network has to become much more efficient (see *Figure 1*).

If, in fact, the whole world is going to be tied together by Internet protocol (IP), one network, then, in fact, these three networks—a voice network, a wireless network, and a data network—have to come together in a single network. These inefficient systems have to become more efficient as we go forward.

Today, in fact, they're tied together to some degree by gateways and services. And increasingly, the people are recognizing that you're not going to have a separate data network; you're not going to have a separate wireless network; you won't have a separate voice network. In fact, all of these networks will be linked together seamlessly, running the same protocol, running the same efficiencies, getting the same volume economics that you get when you have a common infrastructure and a common set of standards.

This is what makes our collective industries work and provide great value to the consumer. We've come up with standards; we've come up with volume economics associated with those standards; and we've delivered those volume economics to the end user.

So, these three networks, which are somewhat glued together today, are going to have to be integrally tied together in the future. And that's what we all refer to as the next-generation network.

We're going to evolve, but we won't throw away the old generations instantly. We all know that—just the same way that we know that when the personal computer (PC) came about, we didn't throw away the mainframe computer. We just built around it. We won't throw away all of the elements of the old-generation networks, either. We will build around them. But we'll have the convergence for the universal IP network. It will be a single network for maximum service, maximum applications, and maximum economics to the end user.

We'll have different ways of accessing this network and the content on that network, and we'll have scalable content delivered to the device at hand.

Now, we've had a lot of press in the last year or two, especially coming out of western Europe, which kind of missed out on the computer revolution but is a leading player in the mobile handset. The press out of Europe basically said that the PC is dead—that everything is going to be mobile handset, the mobile Internet, and mobile commerce. Everything is going to be e-mobile.

FIGURE 1

Today's Network: Inefficient Systems Become Efficient

I think that misses the point entirely. There are various ways to access the Internet, but there will be only one Internet. That one Internet will be the place where content resides. If you access that content with your handheld device or your PDA, you'll get scaled content. If you access it via your PC, you'll get the same content at a different scale. It will be essentially the same content, just scaled for delivery.

Just as our computing industry for years has looked at evolutions of technology, we're hoping to see the same thing in the bandwidth area. Look at the line-up of Intel processors going back to the first 8-bit and 16-bit processors that populated the first IBM PC, which is today just 20 years old, and that were running basically a megahertz or less when they were introduced; then look at today's state-of-the-art processors, which run nearly at 2 gigahertz (GHz). You can see at what sort of speed that transition has occurred.

We joke in our industry about Moore's law, which says processing speed doubles every 12 or 18 months, while in the telecommunications industry, bandwidth doubles every century. So the two industries are integrally tied together here. They're just a little bit out of step.

When it was first introduced, the Internet speed was, in fact, 300 baud; so we have come up a little bit from that. But there is a decided shortage of high-bandwidth content capability to the users, especially to small/medium enterprises and the home consumer in the United States today. Those consumers want more bandwidth. The computers that they have can handle a lot more bandwidth. And if you want rich content going into the home, whether that content is information, entertainment, or commerce-related content, they need more bandwidth to achieve that.

And not just the United States, clearly. You need that content around the world. Some countries are a little more progressive in this space than we are. Korea, for example, I think has the highest density of high-bandwidth capability to the home of any country in the world. And that's a result of the government fostering that installation plus strong competition between digital subscriber line (DSL) and cable in Korea.

We have yet to really establish that strong competitive spirit here in the United States to achieve that same sort of penetration that Korea or Singapore or other countries have. It's estimated that in the United States, only about half of the people who really want high-bandwidth capability to their home will get it this year. Two million or so DSL users and a similar number or so of cable users are faced with increased difficulty to get that content, increased difficulty in servicing that content, and most recently, the realization that the providers are, perhaps, more interested in profit margins associated with delivering that than they are with expending the capital outlay and expanding the service to their users. So they're starting to raise prices.

If you want to look at what volume economics can bring, though, by a single-standard capability, the transistor world is a good example. *Figure 2* shows the number of transistors that are produced each year by our industry worldwide: something like 200 quintillion transistors per year. That is roughly equal to every printed character, letter, or number that has been recorded in the history of mankind. So each year we make the same number of transistors as every character that man has written down in man's lifetime.

You can see that this number continues to grow at a logarithmic rate in terms of the increase and that the average price per transistor continues to decrease the logarithmic rate. That's why you get the anti-inflationary nature of our industry in terms of bringing more capability for the same cost, year upon year upon year.

Metcalf's law really drives the networking industry, and it basically says that if you double the number of intersections

FIGURE 2

Transistors Shipped versus Average Transistor Price

to a network or network and patch points, then the network becomes more valuable. And if you have more and more contact points, have interoperability between networks, and have competition, then you start to get the same thing in the network space that we've had in the transistor space for the last 30 years.

Figure 3 shows the cost per unit of bandwidth for different styles of bandwidth, different communications protocols. There's optical carrier (OC)–3 on up to OC–192, and then at the bottom there's 10 Gigabit Ethernet. And 10 Gigabit Ethernet happens to be a longstanding standard protocol that's been used in local-area networks (LANs) for years and is available almost for free in terms of bandwidth capability in the LAN.

It seems pretty clear that the LAN standard and the long-haul or metro network standard are going to merge, probably in OC–192 and 10 Gigabit Ethernet. They're both at the same bandwidth. There's no reason to have two separate

FIGURE 3

ASP per Gigabit of Bandwidth

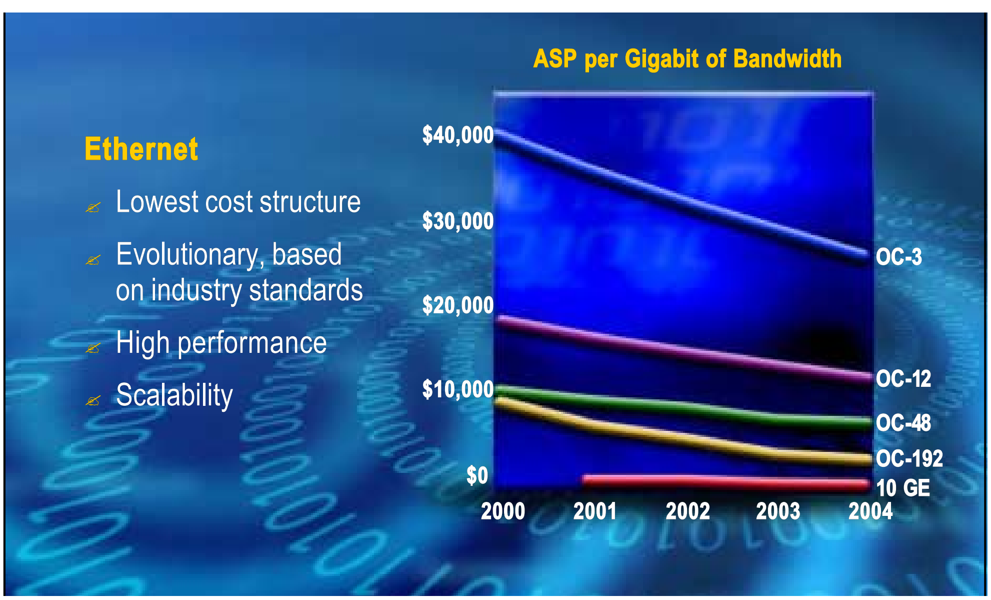

protocols, per se. If you had one, then you can do without a lot of the intermediate switching, redirection, and gateway gear. It's directly traveled.

But the issue here is that open standards beget volume economics, beget value to the end user. Very, very low cost, similar to the economics of transistors. And "open standards" here means that the point of origin and the point of destination adhere to the same protocols. And whether we're talking about the application stacks, the transport stacks, the networking capability, the data links, or the physical interface devices, once we get to common standards, then, in fact, we can start to get volume economics.

And the real beauty of this sort of open-standard capability is it really opens up the arena of competition—most specifically for applications. We can sit and talk about the merits and technical details of Ethernet versus OC–192, and the end consumer just really doesn't care. The end consumer is interested in the solution, the capability: What do you bring me? What can I do with it? What enjoyment do I get? What's the user interface look like? As soon as you start to have open standards, you open up the entire industry to people who can bring applications and port those applications from machine to machine, locale to locale. Then, in fact, you start to increase the rate of innovation and bring the end consumer more value.

So this open interface concept, fostering competition, fostering the rate of technology moving forward is very, very key to the whole message today. This is what's happening with the convergence of computing and communications.

Industries that historically were built up with rather stovepipe proprietary standards are being opened up and are allowing the whole industry to innovate around them, rather than just one company innovating around a particular proprietary standard.

So I think that this next-generation network—this collision bringing the best computing standards together with the best communications standards, and then allowing the whole industry to innovate around those standards—will keep technology moving forward very rapidly.

This paper was excerpted from the author's keynote speech at Networld + Interop in Las Vegas, May 2001.

Optical Networking: Recent Advances, Trends, and Issues

Naveen Bisht

Chief Executive Officer, President, and Founder

Nayna Networks, Inc.

Raj Jain

Chief Technology Officer and Co-Founder

Nayna Networks, Inc.

Abstract

This paper summarizes the latest optical telecom networking market developments. Issues that are being debated hotly in the networking community are presented. This includes the issue of traffic glut versus growth, optical versus electrical switching, Ethernet versus synchronous optical network (SONET), and mesh versus ring. We then also explain the trend toward optical cross-connects.

Introduction

The telecommunications industry experienced a rapid downturn in 2001. As a number of industry gurus attempted to "call the bottom," some of the most generally accepted beliefs turned into mirages. Take a look at the metro market. It was believed that the metro market was poised for strong growth. Institutional money flowed into this sector and created dozens of start-up competitors in every segment, each claiming its own position and striving to provide cheaper, faster, and smaller products. However, the demand turned out to be smaller than was anticipated. On top of that, with the obvious price competition that results in such a situation, a seemingly great opportunity now looks significantly gloomier. The question, which comes to the mind of every telco professional, is "have we hit the bottom yet?" Perhaps not. There are a few factors that will have a large impact in 2002 and ahead:

1. Capital expenditures budgets for carriers are still declining by about 24 percent in 2002.
2. Carrier spending is transitioning from large network build-outs to incremental bandwidth expansion based on immediate revenue-generation models and through the extraction of more efficiency out of the deployed capacity.
3. Oversupply of optical components caused by consolidation and bankruptcies of some of the carriers are expected to last through at least the first half of 2002.
4. Spending disruptions due to carrier consolidation, management changes, and reorganizations, etc.
5. Pricing pressures at all levels of the food chain from systems and components to services.

A silver lining, however, exists among this gloom. Optical systems based on new technologies are going to offer new services opportunistically at a reasonable cost by incorporating new functionalities. Much of the innovation is required to enable carriers to offer new, higher-margin services and to reduce their capital and operational expenditures and network costs. These will be accomplished by the innovation at the component level, at the systems level by adding these innovative components with intelligent software, and at the architectural level by leveraging new architectures that result in cost savings and higher margins.

Ultimately, the end-user customer, whether it's business or residential, will drive growth in our industry. The days of "get to market fast, build it, and customers will come with open arms" are gone, and the mantra going forward is success-based, revenue-based, and incremental investment in the network. On the business side, lambda services are gaining traction with the carriers, and these services aim to provide line functionality at high speeds and competitive costs. Although 2002 may still not look very good for us in the industry, the foundation for innovation in the optical networking industry has been set. This innovation offers a compelling value proposition for the carriers to use as they start planning their next phase of network upgrade, build-out, and the maximizing of the efficiency of their networks.

In spite of the economic downturn, it is important to note that the number of Internet users has grown consistently during the past five to six years. In the late 1990s, UUNET traffic was almost doubling every 90 days. While the current rate of growth may not be as high as during those times, it is still growing at a rapid pace. This growth will provide opportunity for the system vendors and component vendors to be able to leverage their innovative products to build into systems for the carrier infrastructure upgrade that they will have to do in order to remain competitive in the market place.

Each technology follows an S-shaped curve, as shown in *Figure 1(a)*. If we were to plot the number of problems solved as a function of time required to solve them using a given technology, the curve would have an S shape. In the beginning, when a technology is just evolving, it takes quite a bit of time to solve very few problems. In financial terms, it means that the technology developers require substantial funds but that the resulting commercial value may be little. This phase of the technology is called the "research" phase and is funded by governmental research-funding organizations such as the National Science Foundation (NSF) or the Defense Advanced Research Project Agency (DARPA) in the United States. After the fundamental problems have been solved, the technology growth curve takes an upturn, and it takes very little time to solve many problems. In other words, it takes very little money to produce quite a bit of money using that technology. The technology is then taken over by the commercial sector. Finally, after all easy problems have been solved, the technology growth curve takes another turn. Now only hard problems are left and it takes quite a bit of time to solve very few problems. In other words, the technology is no longer cost-effective. At this point, the researchers and the commercial world move on to some newer technology.

Computer networking has also followed an S-shaped curve. If we were to plot either the number of hosts on the Internet or the total number of Internet users, we would find that the curve is S shaped.

The technology started with the concept of packet switching in 1968 and the initial implementation of four-node ARPAnet in 1969. The knee of the networking growth curve occurred in 1995 with the popularity of the World Wide Web (WWW). The peak growth rate (the center point of the S curve) happened sometime in late 1999 or early 2000. At this point, we are slightly beyond the center point as shown in *Figure 1(b)*. This is true for the electronic networks. Conversely, optical networking is still near the knee. There is quite a bit of research activity going on and the commercialization of the optical networking technologies is just starting to take place. We should soon see an exponential growth in optical networking.

There are four issues that are being hotly debated this year at networking conferences. These are as follows:

1. Bandwidth Glut versus Traffic Growth
2. O–O–O versus O–E–O
3. Ethernet versus SONET
4. Mesh versus Ring

We explain both sides of these debates next in this paper and also explain the trend toward optical cross-connects.

1. Bandwidth Glut versus Traffic Growth

One of the fundamental issues that are being debated is whether network traffic growth has stopped. The starting point of this debate was a forecast by McKinsey & Co and J.P. Morgan in May 2001 that Internet traffic growth will slow down from 200 to 300 percent per year to 60 percent by 2005. Soon after that, *The Wall Street Journal*, *The New York Times*, *Forbes*, and other popular media papers and magazines reported the finding that 98 percent of the fiber is unlit. Nortel Networks blamed the loss of its revenue on the falling Internet protocol (IP) traffic. Michael Ching, a Merrill Lynch analyst, reported that carriers are using only average 2.7 percent of their total lit fiber capacity.

All of these gloomy statements and forecasts accentuated the fall of the telecom market. However, when one analyzes

FIGURE 1

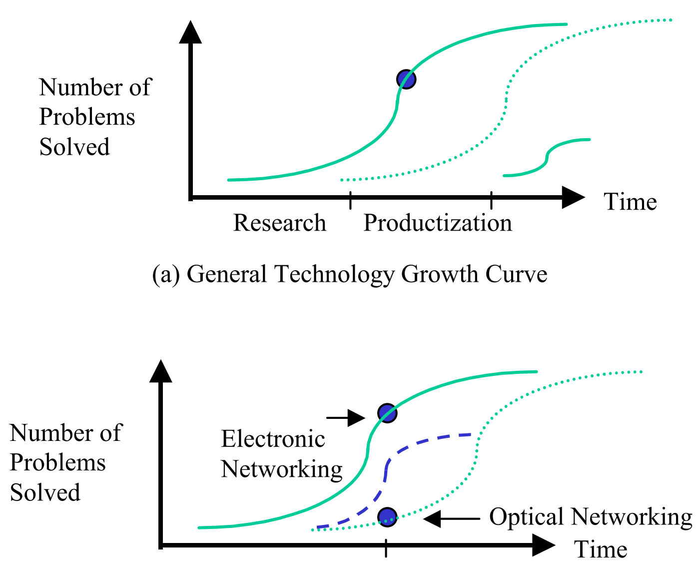

(a) General Technology Growth Curve

(b) Electronic versus Optical Networking Growth these statements, one finds that while these statements might be true, they should not be a cause for alarm. Of the three aforementioned statements, the first one is a forecast while the other two are based on current usage. While the forecast may or may not be true, the current usage needs to be explored further. Installed fiber alone does not constitute the telecom infrastructure. Actually, fiber installation costs are high compared to the costs of the fiber itself, and so it is quite common for carriers to install cables with hundreds of fibers when they need just a few fibers. Having unlit fiber is quite common and should not be considered alarming. Similarly, the average usage of networking facilities is generally very low. For example, most of the computers today are equipped with 100 megabits per second (Mbps) Ethernet ports. The average usage of these ports is very low—generally less than 1 percent. Nevertheless, networking links are often the bottlenecks, and we find ourselves waiting for the information to arrive over the network. At those peak usage times, we need all 100 Mbps and more. Thus, the network capacity is planned for peak usage and not the average.

Telechoice, a market research company, conducted a usage study for Williams Communications. They studied the 22 most used routes in the United States and found that the utilization on 14 of these routes exceeded 70 percent. This high utilization would necessitate upgrading these routes and the networking equipment. Larry Roberts, one of the early founders of the Internet, himself conducted a measurement study of traffic on 19 of the largest Internet service providers (ISPs) in the United States. He obtained data on 95-percentile usage at each of these ISPs at six months intervals between April 2000 and April 2001. A plot of this data is shown in *Figure 2*. The figure also shows 95-percentile of the total traffic on the measurement dates. Roberts concluded that the traffic growth factor was 390 percent per year during the April 2000 to October 2000 and that it was 400 percent during the next six months. This study has been very helpful in calming down the bandwidth glut debate—at least for the time being.

2. O–O–O versus O–E–O

Gordon Moore, one of the founders of Intel, observed that the number of transistors in integrated circuits doubles every 18 months. This observation has been found to be true over the last two decades and is often interpreted to mean that the speed of electronic circuits is doubling every 18 months. Even this high rate of growth is slow when compared to the growth of network traffic. *Figure 3* shows the growth of network traffic and the speed of electronic circuits as a function of time. The traffic growth has been assumed to be 400 percent per year using the aforementioned Roberts study. Note that within a few years, the network traffic growth is expected to be several orders of magnitude greater than what can be handled by electronic circuits without parallelism. This is an argument in favor of optical switching.

In current switches, even though the signal comes in optically on a fiber and goes out optically on a fiber, it must be converted to electronic form for switching. This is called optical-to-electrical-to-optical (O–E–O) switching. In optical-to-optical-to-optical (O–O–O) switches, the signal is switched optically. There are several methods for optical switching, the most common being the use of micro mirrors using a micro-electrical machines (MEMs). Based on the switching instructions, the mirrors can be rearranged quickly to make the light coming on one fiber go to another fiber as desired.

O–O–O switches are data-format independent in the sense that the data being switched could have any data link format, such as SONET or Ethernet or synchronous digital hierarchy (SDH) or optical transport network (OTN) (based on ITU G.709). The same O–O–O hardware can support all data link formats. An O–E–O switch, on the other hand, will require a different circuit or software to support each of these formats.

The O–O–O switches are also relatively rate independent. The same hardware can switch a 2.5 gigabits per second

FIGURE 2

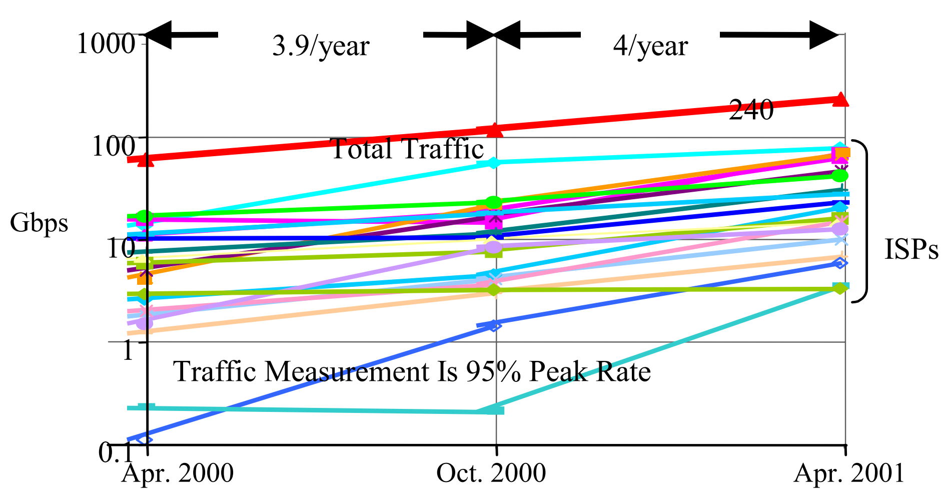

(Gbps) signal or 10 Gbps signal. Of course, the noise tolerances are tighter for 10 Gbps than 2.5 Gbps, and so the component quality has to be better for 10 Gbps equipment than those designed for 2.5 Gbps. O–E–O switches, on the other hand, are highly rate dependent. An integrated circuit designed for 2.5 Gbps cannot handle 10 Gbps. In the simplest approach, four 2.5 Gbps circuits will be required to handle a 10 Gbps signal, and so the cost of O–E–O switches grows proportionally to the data rate. The same argument applies for space and power. This is the prime reason in favor of O–O–O switching at high data rates. The conventional wisdom at this time is that, at 10 Gbps and higher rates, the per-port cost of O–O–O switches is less than that of O–E–O switches.

The rate independence of O–O–O switches also reflects in their upgradability. Upgrading a 2.5 Gbps O–O–O switch to 10 Gbps may require fewer changes than those in an O–E–O switch. As mentioned earlier, for an O–O–O switch, the same basic design can be upgraded with higher quality components to support the higher rates. Upgrading an O–E–O switch, on the other hand, will require a complete change of design.

The next two issues are related to the ability to handle part of a wavelength or multiple wavelengths. An O–E–O switch can easily separate out a part of a signal, and so adding or dropping a part of the wavelength is easily accomplished. In an O–O–O switch, this can be done only if the different parts of the signal are different in some optical characteristics, such as time (slots), frequency, phase, or polarization. On the other hand, an O–O–O switch can switch a multiwavelength signal as easily as a single wavelength, provided that the optical components are properly designed. O–E–O switches will require separate circuits to handle each of the wavelengths and would be very costly for the large number of wavelengths that can be accommodated in a fiber.

In terms of performance monitoring, O–O–O switches are bit-rate– and format-independent and so cannot easily see bit errors (which are basically rate or format errors). O–O–O switches monitor optical defects such as wavelength shifts, optical signal-to-noise ratios, or power levels. These defects also result in bit errors, but not all bit errors are visible to optical monitors. To monitor electrical signals, O–O–O switches provide optional electrical monitoring. In O–E–O switches, the cost of monitoring is already built in, since they have to verify the data rate and format before they can switch the signal.

If two or more signals need to be sent on a single fiber, they should have different optical characteristics (wavelength or polarization). Therefore, it may sometimes be necessary to convert these optical characteristics (primarily the wavelength) at switching points. Optical approaches for wavelength conversion are a few years away. Several companies have announced products in this space, but none is shipping at this point. Therefore, O–O–O switches offer optional O–E–O–based wavelength conversion on selected channels as needed. Often, this is not needed, since the dense wavelength division multiplexing (DWDM) equipment, which is used with both O–E–O and O–O–O switches, already has transponders that change wavelengths as required.

A comparison of O–E–O and O–O–O switches is summarized in *Table 1*.

In summary, O–O–O switches are the direction for the future as we transition into higher speeds and larger switch capacities.

3. Ethernet versus SONET

It is now well established that the traffic on most carrier networks is predominantly data traffic. SONET/SDH technology is designed for voice traffic and is very expensive compared to Ethernet, which is designed for data. It is clear that using Ethernet switches in place of SONET add/drop multiplexers (ADMs) will reduce the cost considerably. However, there are several obstacles to the adoption of the Ethernet technology, the primary ones being the reliability and availability.

FIGURE 3

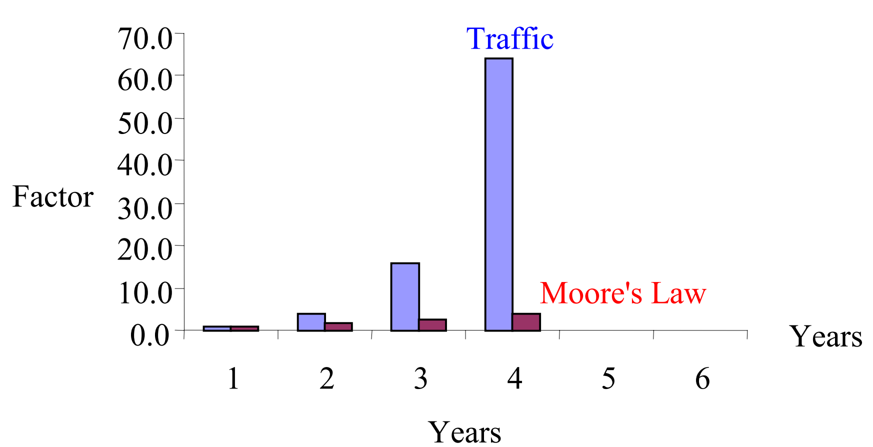

TABLE 1

O–E–O versus O–O–O

Feature	*O–E–O*	*O–O–O*
Data Format Dependence	Yes	No
Cost/Space/Power Independent of Rate	No	Yes
Upgradability to Higher Rate	No	Yes
Sub-Wavelength Switching	Yes	Future
Waveband Switching	No	Yes
Performance Monitoring	Bit-Error Rate	Optical Signal Degradation
Wavelength Conversion	Built-In	Currently Electronic

SONET technology was designed primarily for carrier networks and has very robust reliability and availability mechanisms built in. In particular, SONET networks are designed to provide 99.999 percent (five nines) availability, which is equivalent to a down time of five minutes per year. This is achieved by a high-level of redundancy inside and outside the equipment. Ethernet technology, on the other hand, was designed primarily for enterprise networks, where availability requirements are not that high.

Carriers' recent interest in Ethernet is visible from the activities in the 10-Gbps standardization efforts. 10-Gigabit Ethernet is designed for two data rates: 10 Gbps for local-area network (LAN) applications and 9.5 Gbps for wide-area network (WAN or telecom) applications. The WAN version uses SONET framing. It is compatible with SONET equipment except for the clock jitter requirements. SONET requires a clock jitter of 4.6 to 20 parts per million (ppm) while 10-Gigabit Ethernet requires only 100 ppm. This decision was highly debated in the Institute of Electrical and Electronics Engineers (IEEE) standards committee and may have resulted in the delay of the standard; however, requiring tighter tolerances would have increased the cost of the equipment significantly. As a result of this, a 10-Gigabit Ethernet signal cannot be sent directly to a legacy SONET ADM. Ethernet line-termination equipment (ELTE) is required to buffer the incoming signal and send out the well-conditioned signal to the SONET equipment. In this way, the extra cost of clock-jitter conditioning is not incurred if the entire WAN is based on 10-Gigabit Ethernet technology. This is the future plan of many carriers, particularly those that are primarily data carriers.

SONET networks are traditionally organized in dual-ring topologies that allow for a very fast recovery from node and link failures. Ethernet equipment is traditionally organized as mesh networks. Ethernet switches use a spanning tree algorithm to automatically convert the mesh topology to a spanning tree topology for forwarding. The spanning tree takes a few minutes to converge, and so the restoration times may be large. The IEEE has started an effort to design a fast spanning tree algorithm.

To provide fast recovery time for Ethernet traffic, the IEEE has also started a resilient packet ring (RPR) project that will allow Ethernet traffic to be sent on dual-ring networks. This will provide fast recovery times matching those of SONET, while at the same time being more efficient in terms of redundancy by allowing both rings to be used when there is no failure.

A comparison of SONET versus Ethernet is summarized in *Table 2*.

Now that the 10-Gigabit Ethernet standards are nearing completion, IEEE has started discussing the next steps. In particular, a survey of IEEE 802.3ae members has indicated that 70 percent of the members would like the next version of the Ethernet to be 40 Gbps rather than 100 Gbps. This will allow the optical carrier (OC)–768 technology being developed for SONET to be reused for Ethernet.

4. Ring versus Mesh

Telecommunications networks are currently organized in ring-based topologies while the data networks use the mesh-based topologies. This has started the debate on the merits of ring and mesh topologies. This debate is similar to that between Ethernet and token rings back in the initial days of IEEE 802 standardization.

As shown in *Figure 4*, on a ring, all links have to have the same data rate. If any one link is upgraded—for example, from 2.5 Gbps to 10 Gbps—all links and nodes have to be upgraded to that rate. Therefore, rings are more suitable for networks where the traffic among the nodes is homogeneous. Rings are not generally used in long-haul networks where the traffic is highly non-homogeneous. The mesh networks, on the other hand, allow incremental upgrades. Any link can be upgraded to a higher rate while the others can remain at a lower rate. Similar arguments apply in the case of DWDM networks, where the number of wavelengths must be the same on all nodes of a ring.

Mesh networks typically require 50 percent less protection and 50 percent less working capacity than rings. This is because of the inherent spatial reuse feature of mesh networks whereby the links not being used by one flow can be used by other flows. The savings from mesh networks increase as the degree of connectivity increases and as the non-homogeneity of the traffic increases.

TABLE 2

Feature	*SONET*	*Ethernet*
Bit Rate (bps)	155M, 622M, 2.5G, 10G, 40G,…	1M, 10M, 100M, 1G, 10G,…
Timing	Isochronous (Periodic 125µs)	Plesio-Isochronous
Multiplexing	Bit	Packet
Clocks	Common	Independent
Clock Jitter	4.6 to 20 ppm	100 ppm (May Change)
Usage	Telecom	Enterprise
Volume	Millions	100s of Millions
Price (10 Gbps)	>10k	<1k
Recovery	50 ms	Few Minutes
Topology	Rings	Mesh

Currently, there are two parallel efforts in the telecom networking community. On one side, they are trying to develop mesh-based protection and restoration mechanisms and protocols, while on the other side they are also trying to develop ring-based mechanisms and protocols suitable for data traffic. This second effort is reflected in the RPR work going on in the IEEE. RPR uses a dual counter rotating ring topology similar to that used in SONET and fiber distributed data interface (FDDI) networks.

5. Optical Cross-Connects

As the network traffic increases, the capacity of cross-connects and switches must also increase. There are two ways to increase the capacity: higher port count or higher speed per port. For example, 80 Gbps traffic can be handled with eight 10G ports or two 40G ports. Increasing the port speed increases the number of multiplexing and demultiplexing (grooming) points and increases the number of hops. This has lead to a trend toward a larger number of ports per switch. A similar argument applies toward the number of wavelengths and the data rate per wavelength as well.

A single cross-connect with a large number of ports is more cost-effective than an interconnection of several smaller cross-connects. To illustrate this, consider the two alternatives shown in *Figure 5* for a 24x24 cross-connect. Alternative A consists of using 24 4x4 switches with a total of 192 ports. Alternative B consists of a single 24x24 switch with a total port count of 48. Since most of the cost of the cross-connect is in the port cost, alternative A is almost four times more expensive than alternative B. In general, larger cross-connects are more cost-effective than multiple smaller ones. This has lead to a trend toward large cross-connects that can be more economically realized with optical technology than electrical.

FIGURE 4

Ring versus Mesh

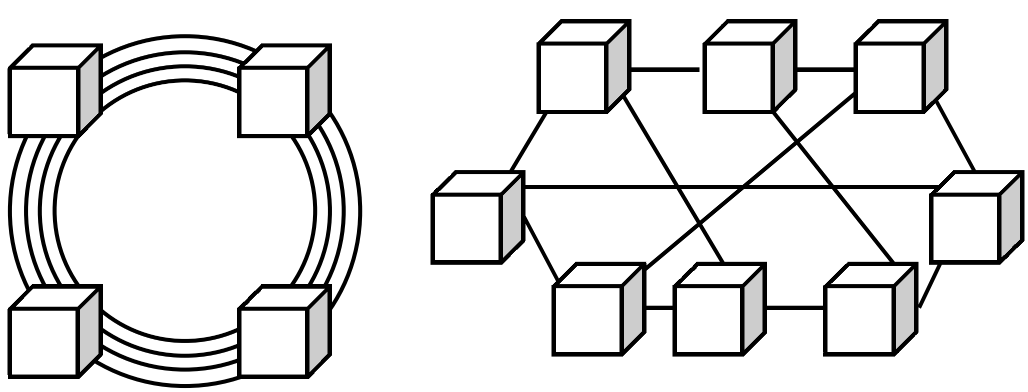

FIGURE 5

SONET versus Ethernet

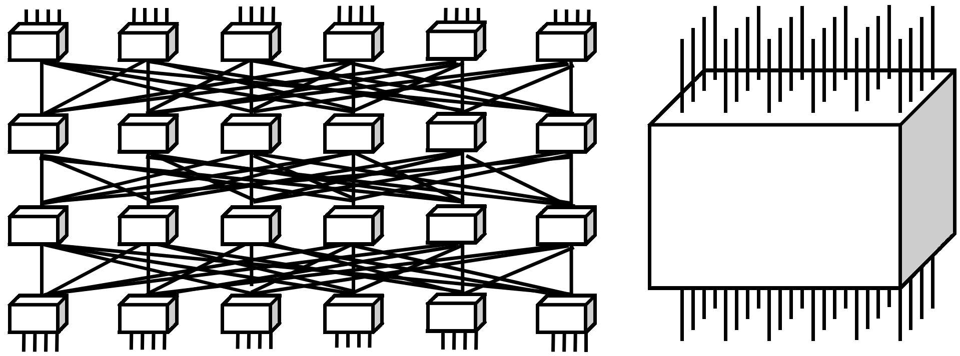

(a) Using Several 4×4 Cross-Connects (b) One Cross-Connect

Summary

The optical networking industry has been affected by the economic slowdown. However, we believe that the slowdown is temporary. Internet traffic is still growing. Carriers need to find ways to increase their revenues from this increasing traffic and to reduce their capital and operating expenditures. Optical cross-connects, larger number of ports, mesh topologies, and Ethernet-based networks seem to be more economical than their counterparts.

Free Space Optics and the Last 500 Meters

Scott Bloom
Chief Technical Officer
AirFiber Inc.

Introduction

There are several key factors that are bringing free space optical technology (FSOT) into the carrier space as a means for broadband access. The first is growing and seemingly insatiable bandwidth demand in the marketplace. FSOT technology provides fiber-like speeds without significant initial capital expenditures for scarce resources such as spectrum. While fiber rings are becoming ubiquitous, getting the bandwidth from the rings to customers remains challenging. While FSOT infrastructure sounds inviting, there are many near-term hurdles for FSOT products. The primary customer concern for FSOT products is availability. This needs to be statistically demonstrated with real FSOT systems in the field before there is widespread customer acceptance of this technology. Other concerns are compatibility with existing local-exchange carrier (LEC) networks and premises networks, cost, carrier-class hardware, ease of installation, and network management.

The physics of the atmosphere dictate some very specific performance limitations on free-space systems that can be derived very generally to put an upper bound on carrier-grade range expectations and claims by many vendors in the industry.

Free Space Optics Subsystems

Figure 1 illustrates the major subsystems in a complete carrier-grade free space optics communications system. The optical apertures on a free space system can have an almost infinite variety of forms and some variety of function. They can be refractive, reflective, diffractive, or combinations of these. In the figure, the transmit, receive, and tracking telescopes are illustrated as separate optical apertures; there are several other configurations possible where, for example, a single optic performs all three functions thereby saving cost, weight, and size. The important aspects of the optical system on the transmit side are the size, which determines the maximum eye-safe laser flux permitted out of the aperture and may also prevent blockages due to birds, as well as the quality of the system, which, along with the f-number and wavelength, will determine the minimum divergence obtainable with the system. On the receive side, the most important aspects are the aperture size, which determines the amount of light collected on the receiver, and the f-number, which determines the detector's field of view. The tracking system optics must have a wide enough field of view to acquire and maintain link integrity for a given detector and tracking control system.

Diode lasers are driven with a DC bias current to put the devices above threshold and are then modulated with an AC current on top of that to provide, for example, on-off keying (OOK) for data transmission. For lasers with output powers below about 50 mW, off-the-shelf current bias and drive chips are available; for higher power lasers, custom circuits or radio frequency (RF) amplifiers are generally used.

The receive detector is coupled to the receive aperture through either free space or a fiber. Depending on the data rate and optical design alignment tolerances can be extremely restrictive. For example, for data rates to 1.25 gigabits per second (Gbps), detectors with relatively large active areas (500 microns diameter) can be used, making alignment to the receive aperture fairly straightforward. For fiber coupling into multimode fibers, the core size is about 63 microns in diameter, and this makes alignment much tougher. The use of special materials or controls is required, however, coupling is more modular.

Detectors are generally either photo intrinsic diodes (PINs) or avalanche photodiodes (APDs). For carrier-class free space optics systems, an APD is always advantageous, since atmospheric-induced losses can reduce received signals to very low levels where electronics noise dominates the signal-to-noise ratio (SNR). Of course the APD has to be capable of meeting the bandwidth of the system. Usually, a trans-impedance amplifier is used after the detector, and in most cases they provide the highest gain at the fastest speed.

Tracking detectors are relatively large-area devices if they are CCD, CMOS, or quad cell detectors and are therefore easy to align to the tracking optics, although care must be taken in manufacture to co-align these optics with the transmit and receive optical axes. For building-mounted free space optical systems, the tracking bandwidth can be very low, sub hertz, since the bulk of building motion is due to uneven thermal loading of the building. These effects occur on time scales of hours. For systems that are to be mounted on towers or tall poles, the tracking bandwidth should be higher, most likely on the order of several hertz at least, to remove wind-induced vibrations.

Acquisition systems can be as crude as aligning a gun-sight to very sophisticated global positioning systems (GPS)–based, high-accuracy, fully automated systems. The choice of this subsystem really depends on the application and number of devices to be put into a network.

The Free Space Optic Link Equation

The link equation for a free space optical system is actually very simple at a high level (leaving out optical efficiencies, detector noises, etc.). The equation is illustrated in *Figure 2*. The amount of received power is proportional to the amount of power transmitted and the area of the collection aperture. It is inversely proportional to the square of the beam divergence and the square of link range. It is also inversely proportional to the exponential of the product of the atmospheric attenuation coefficient (in units of 1/distance) times the link range.

Looking at this equation, the variables that can be controlled are the transmit power, the receive aperture size, the beam divergence, and the range of the link. The atmospheric attenuation coefficient is uncontrollable in an outdoor environment and is roughly independent of wavelength in heavy attenuation conditions. Unfortunately, the received power is exponentially dependent on the product of the atmospheric attenuation coefficient and the range, in real atmospheric situations, for carrier-class products (i.e., availabilities at 99.9 percent or better). This term overwhelms everything else in the equation. This means that a system designer can choose to use huge transmit laser powers, design large apertures, and employ very tight beam divergences, and the amount of received power will remain essentially unchanged. *Figure 3* graphically illustrates this point. Except for clear air, the atmospheric loss component of the link equation dominates by many orders of magnitude, essentially washing out any system design choices that could affect availability. For free space systems that are designed to be carrier-class, this fact needs to be accepted and the system design produced with this in mind. Basically, the only other variable under the designer's control is link range, which must be kept short enough such that the atmospheric attenuation is not the dominant term in the link equation. As will be discussed in a later section, this implies that the link range *must* be kept at less than 500 meters. If this is realized, then efficient designs can be produced that permit economical, reliable operation under this constraint.

Figure 4 shows a tabular link budget for the top five free space optical system manufacturers. Most of the system parameters are freely available from manufacturers data sheets; where a parameter has been assumed, it is noted. One of the vendors (highlighted) has a product at 155 megabits per second (Mbps) and really should not be compared with the others at 1.25 Gbps but has been included for informational purposes. The atmospheric attenuation condition is 200 decibels (dB)/km. The link ranges were adjusted for each system such that the margin came out to approximately zero. It is interesting to note the wide variety of aperture sizes, wavelengths, transmit divergences, and transmit powers (all of the adjustable system parameters) that have been employed in these systems, yet as previously discussed, the maximum link ranges are all about the same (+/– 30 meters or so), again illustrating the point that there is not a lot that system designers can do to increase link

FIGURE 1

Free Space Optics: Major Subsystems

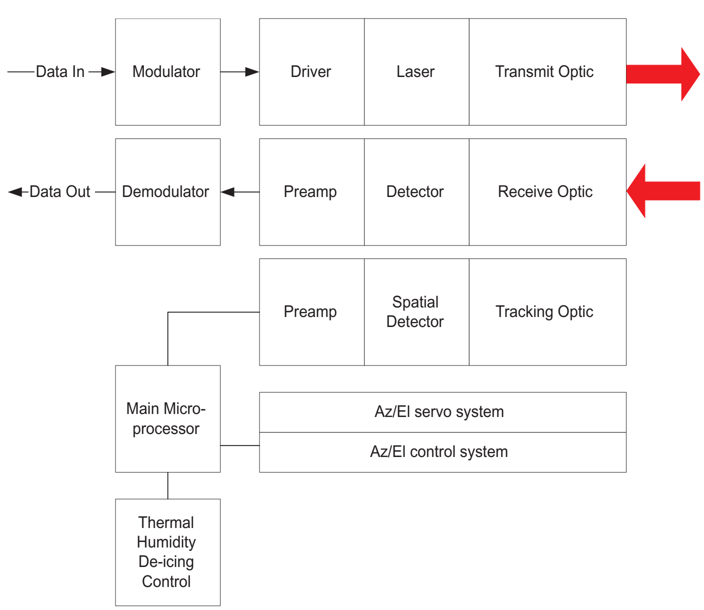

FIGURE 2

Basic FSO Link Equation

Receiver area

$$P_{received} = P_{transmit} \cdot \frac{A_{receiver}}{(Div \cdot Range)^2} \cdot \exp(-\alpha \cdot Range)$$

Beam divergence

Atmospheric Attenuation Factor (dB/km)

range in realistic carrier-grade atmospheric conditions. Given that all manufacturers have about the same performance in high attenuation conditions, *Figure 5* plots the maximum range versus atmospheric attenuation (dB/km) for one of the vendors, which is about the same for all of the vendors. It is interesting to note that at longer ranges (about 1 km or so) the maximum attenuation allowable is about 30–40 dB/km. It turns out that heavy rain can produce this range of attenuation values, and therefore the long link range issue is going to be widespread throughout the world, and not limited to just foggy coastal areas.

The key issue for carriers in deploying free space optical systems is the system availability. The system availability has many factors to it, including equipment reliability and network design (redundancy for example), but these are well known and fairly quantifiable. The biggest unknown is the statistics of atmospheric attenuation. While almost all major airports around the world keep visibility statistics (which can be converted to attenuation coefficients), the spatial scale of visibility measurements is rough (generally 100 meters or so), and the temporal scale is infrequent (hourly in most cases). With the spatial and temporal scales being crude, estimates of availability for carrier grade equipment (99.9 percent or better) are going to be limited to 99.9 percent or worse. These huge databases are therefore not useful except for estimating the lowest acceptable carrier grade of service. To permit carriers to write reasonable service-level agreements (SLAs), better data must be taken. At AirFiber, we have had instruments capable of acquiring this data running continuously for roughly the last couple of years. These instruments, which include a nephelometer and an in-house designed and built weather benchmark system, provide data at the correct spatial and temporal resolution for accurate estimates of availability versus link range.

Figure 6 shows a plot of this data for two cities, Tokyo and San Diego. The one line shows the cumulative probability

FIGURE 3

Geometric and Atmospheric Loss Components

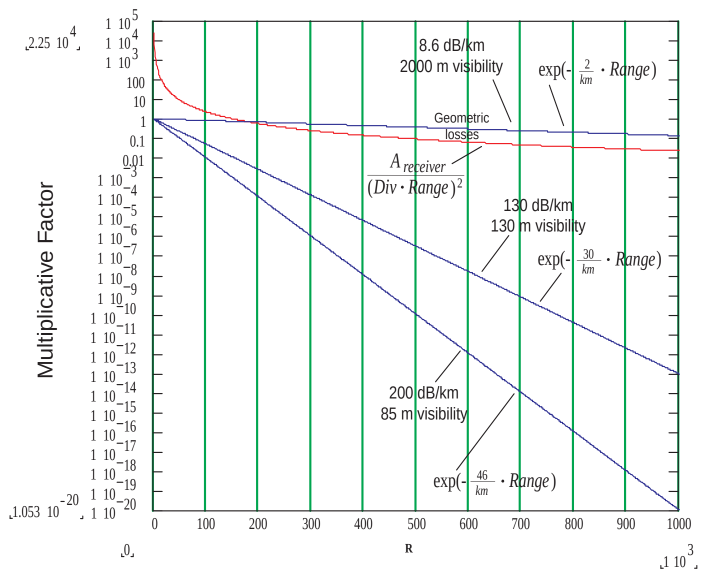

FIGURE 4

FSO Link Range Performance Comparison

	AirFiber	Vendor A	Vendor B	Vendor C	Vendor D	
AtmosLoss						
dB/km	-200	-200	-200	-200	-200	
Transmitter	AlGaAs					
Modulation Format	NRZ OOK					
Receiver	Si APD					
Wavelength	785	850	1550	1550	850	nm
Range	200	190	215	205	176	m
Data Rate	1250	1250	1250	155	1250	Mbit/s
Average Laser Power	18	30	1000	320	15.6	mW
Peak Laser Power	36	60	2000	640	31	mW
Transmit Aperture	5	5	16	2.5	5	cm
Transmit Divergence	0.5	2	2	4.25	2	mrad(1/e^2)
Receive Aperture	7.5	20	40	20	19	cm
Optical Background	0.2	0.2	0.2	0.2	0.2	W/m^2/nm/sr
Receiver FOV	3.25	3.25	3.25	3.25	3.25	mrad(1/e^2)
Receive Filter Width	25	25	25	25	25	nm
Receiver Sensitivity	800	1000	1000	1000	1000	nW
BER	1.00E-12	1.00E-12	1.00E-12	1.00E-12	1.00E-12	
Peak Laser Transmit Power	-14.43697499	-12.2184875	3.010299957	-1.93820026	-15.08638306	dBW
Extinction Ratio Degradation	-0.2	-0.2	-0.2	-0.2	-0.2	dB
Transmit Optics Degradation	0	0	-15	0	0	dBW
Pointing Loss	-1	-1	-1	-1	-1	dB
Geometric Range Loss	-2.498774732	-5.575072019	-0.628169285	-12.78225591	-5.355781251	dB
Atmospheric Loss	-40	-38	-43	-41	-35.2	dB
Atmospheric Scintillation Fade	-1	-1	-1	-1	-1	dB
Receive Optics Attenuation	-1.4	-1.4	-1.4	-1.4	-1.4	dB
Bandpass Filter Loss	-0.7	-0.7	-0.7	-0.7	-0.7	dB
Misc Loss Elements	0	0	0	0	0	dB
Received Peak Power at Detector	-61.23574972	-60.09355952	-59.91786933	-60.02045617	-59.94216431	dBW
Required Peak Power at Detector	-60.96910013	-60	-60	-60	-60	dBW
Link Margin at Range	-0.266649594	-0.093559515	0.082130672	-0.020456169	0.057835688	dB

density function for San Diego and another line shows the same for Tokyo. Also plotted is the link budget equation for an AirFiber product; other vendors' products will have about the same link margins. The left vertical axis shows the percentage of time that the attenuation is greater than or equal to a given value. The horizontal axis is attenuation in dB/km, and the right vertical axis is the maximum link range at zero link margin. To use the chart, one must choose an availability, say 99.9 percent (as shown by the dotted line), move horizontally to the desired city (say Tokyo in this example), move vertically to the link budget equation, and finally move horizontally to the maximum link range, (in the case of Tokyo is about 350 meters). It is interesting to note that Tokyo is qualitatively in the top 10 percent of cities for clarity of the atmosphere and that San Diego is in the bottom 10 percent. Therefore, for most deployments, the maximum range will fall somewhere in between these two cities, certainly less than 500 meters in most cases.

FIGURE 5

Link Range versus Attenuation for a Typical FSO Vendor

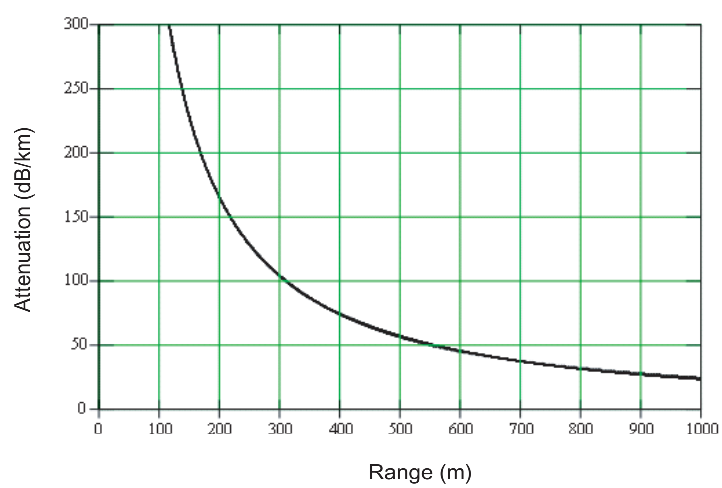

Theoretical preliminary link budget at 1.25 Gbit/s, 10^{-12} BER

FIGURE 6

FSO Availability to 99.999 Percent Resolution for Tokyo and San Diego

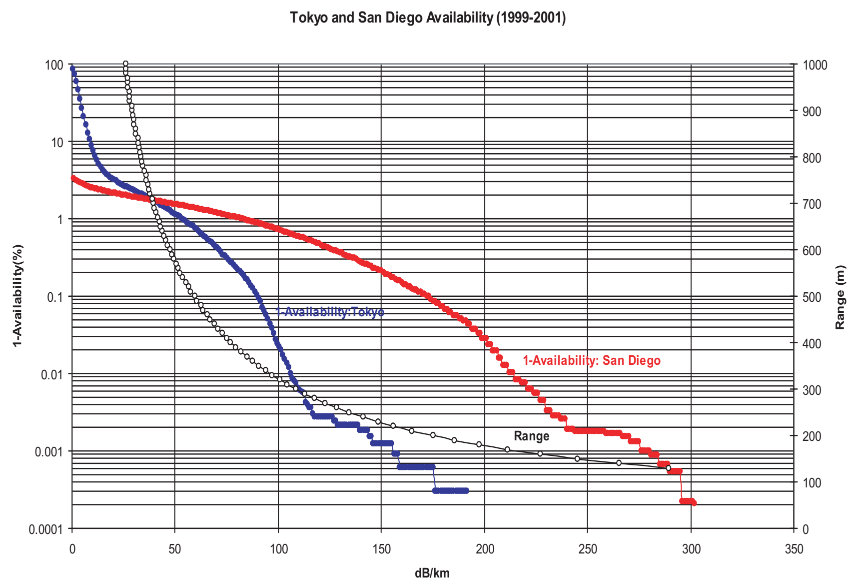

As a final statement about the performance of free space optical systems in realistic weather conditions, we perform some calculations of maximum link ranges for a straw man theoretically "perfect" but as yet unobtainable—the free space optical system. *Figure 7* illustrates the assumptions and results of this calculation. Under the assumptions, it is worthy to note that the detector, electronics, and background noise were set to zero and that the collection aperture was assumed to receive *all* transmitted photons—that is, beam divergence was set at zero. In addition, a very large transmit aperture was chosen so that 80 watts of transmit power could be used and still maintain eye safety. Running through the numbers for 200 dB/km and 10^{-12} bit-error rate (BER) at 1 Gbps data rate, we arrive at a maximum range of about 500 meters. So, even in an "unobtainium" system, the maximum range in realistic atmospheric attenuation situations is about 500 meters.

Another topic that frequently comes up concerning the performance of free space optical systems is the issue of atmospheric propagation and wavelength. One generally held belief is that systems operating at longer wavelengths have better range performance than systems at shorter wavelengths. *Figure 8* shows several calculations performed using MODTRAN for a 1 km path length for which the x-axis is wavelength. The conditions were a visibility range of 200 meters, typical of an advection fog. The y-axis in the top panel is transmission from a minimum of 0 to a maximum of 1. The top panel shows the amount of absorption due to water only in the atmosphere. Here, there are many wavelengths that propagate very poorly due to absorption by water vapor, particularly near 1.3–1.4 microns. The next panel down shows absorption due to oxygen and carbon dioxide, which are relatively narrow lines and are easily avoided.

FIGURE 7

Straw-Man Range Calculation for a "Perfect" FSO Link

Assumptions

– Eye-safe laser
– No detector/electronics or background noise
– **No geometrical loss**
– 1 Gbit/s
– 25.4 cm aperture
– 80 W transmit power
– 200 dB/km Attenuation
– Need SNR=14 for 10^{-12} BER

Results

$$SNR = \frac{1}{2}\,\frac{\eta\, P_{signal}}{h\nu B}$$

$$80W\;10^{-\frac{200\,Range}{10}} = 8.9nW$$

Range = 497 meters

FIGURE 8

MODTRAN Calculation of 785 nm versus 1550 nm Propagation

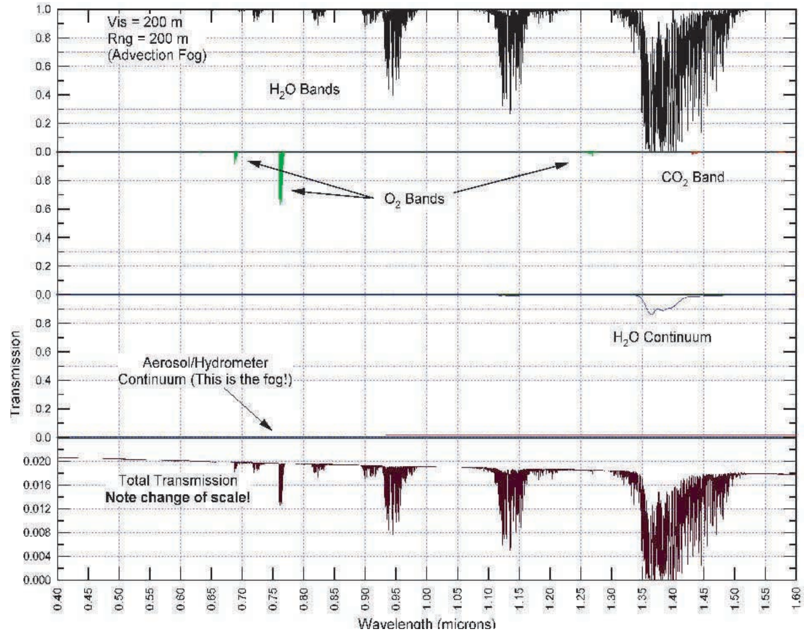

The third panel down shows the effects of Mie scattering by water droplets in the fog. Clearly, this is the dominant loss mechanism under these conditions and is basically independent of wavelength (it is actually a little *worse* at 1.5 microns than at 785 nm, for example). Finally, the bottom panel shows the combined effects of all three loss mechanisms. Again, the result is basically independent of wavelength. There is no advantage in propagation range by using longer wavelengths. Taking this one step further, the calculations were performed at even longer wavelengths, as illustrated in *Figure 9*. Here, the x-axis is again wavelength, the y-axis is the combined loss over a 1 km path. In these illustrations, the vertical axis is attenuation in dB/km. For all wavelengths up to about 12 microns, the result is the same: there are no "spectral holes" in which it is very advantageous to propagate—the attenuation is basically flat over these wavelengths. Finally, the same calculations were carried out all the way to millimeter waves,

FIGURE 9

MODTRAN Calculation out to 10 Microns

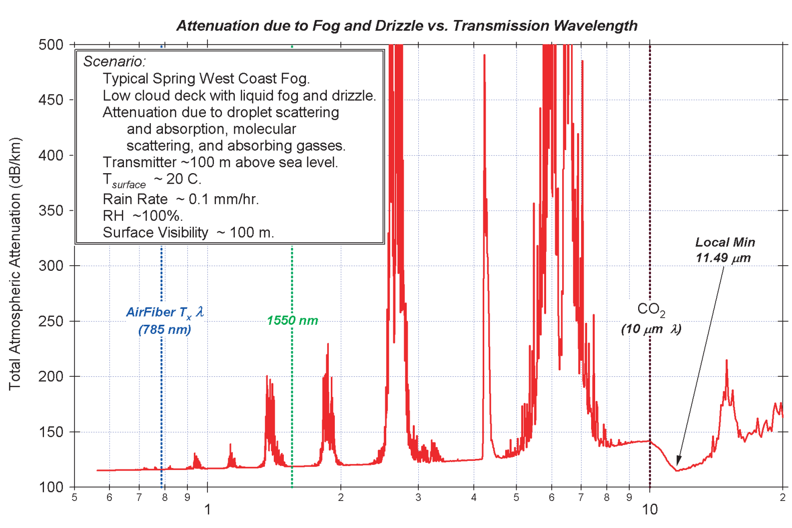

as illustrated in *Figure 10*. This was done for completeness and to make sure the attenuation reduced at RF frequencies to generally accepted values. Not until the wavelength gets to sub-millimeter size (RF) does the attenuation get markedly reduced.

In summary, we can clearly state that for the majority of cities around the world, the carrier-class distance (as defined by 99.9 percent availability or better) is less than 500 meters. In addition, despite numerous claims, all free space optics vendors have about the same range performance in carrier-grade conditions (99.9 percent or better) due to complete domination of the link budget equation by the atmospheric attenuation factor in high attenuation situations. Finally, wavelength has absolutely no effect on propagation range under carrier-grade conditions for wavelengths from visible all the way up to millimeter wave (RF) scales.

Free Space Optics Systems Enhancements

One issue that is debated among free space optics vendors is the necessity of having a tracking system. That is, does the hardware need to compensate for small amplitude, low frequency, and building motion, usually thermally induced, in order to maintain link integrity, or is it sufficient to spread the divergence of the transmitter beam large enough to encompass all foreseeable building movement amplitudes? Going back to the link equation, it is easy to see that spreading the beam divergence impacts the link margin by the inverse square of the spread. In other words, every doubling of the beam divergence negatively impacts the link margin by 6 dB. In addition, even if the beam is spread broadly enough to accommodate building motion, the overall system link margin will fluctuate negatively and/or degrade, since precise alignment of the free space optical systems, by design, will not always be achieved over time and temperature.

Figure 11 illustrates this effect with actual data from a fielded system (distance of 500 meters) running live customer traffic in Phoenix, Arizona. The upper graph shows the change in azimuthal position as a function of time. The x-axis units are arbitrary, however, the pattern of peaks and valleys is due to diurnal thermal changes at the site. The same is true of the bottom graph but for the elevation axis. One tick on either vertical axis corresponds to about 100 microradians of motion, so the maximum azimuthal variation (building twist) was about 1.5 milliradians, and the maximum elevation variation was about 2.5 milliradians. In fact, we have observed elevation variations of 8 milliradians on this same link under different thermal conditions. Another interesting point to note is that a competitor's non-tracking units were removed from this link because they consistently lost link with diurnal thermal cycles. A calculation of the lost link margin for non-tracking units due to building motion is now in order. If the $1/e^2$ beam divergence is 2 milliradians and the links are perfectly aligned at installation, then a motion of 1 milliradian would incur a loss of link margin of about 8.6 dB. Even if the link stayed up during this excursion, the system performance would be degraded by 8.6 dB from the initially specified margin. For a typical carrier-grade link of 200 meters, this means that not having tracking means the system can withstand about 43 dB/km *less* atmospheric attenuation than the initial design. For calibration, heavy rain can induce about 17–40 dB/km of attenuation, so a system that might be designed to handle heavy rain could now be down in a severe storm.

Atmospheric scintillation, as used in this document, can be thought of as changing light intensities in time and space at the plane of the receiver detecting the signal from a transmitter at a distance. The received signal level at the detector fluctuates due to thermally induced changes in the index of refraction of the air along the transmit path. The index changes cause the atmosphere to act like a collection of

FIGURE 10

MODTRAN Calculation to Millimeter Waves

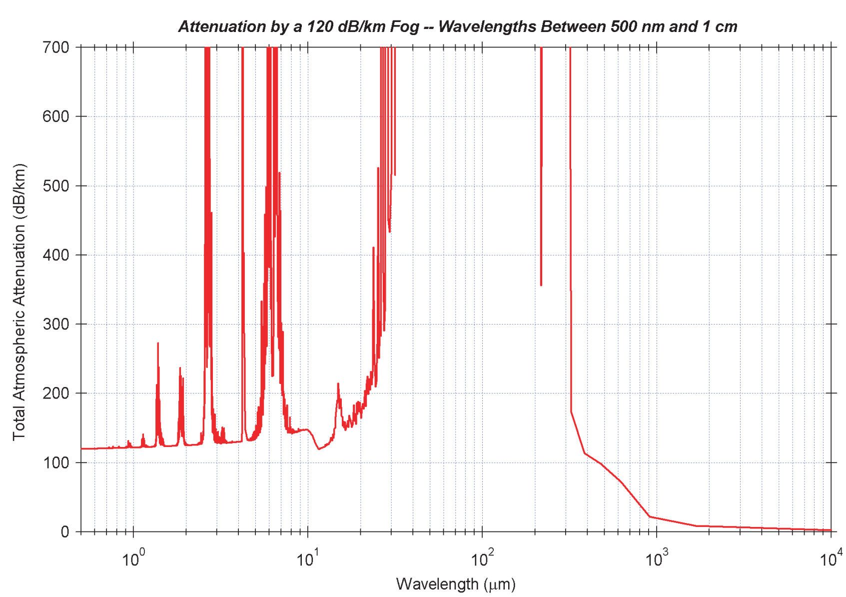

FIGURE 11

Tracking Data—Live Customer Network in Phoenix, AZ

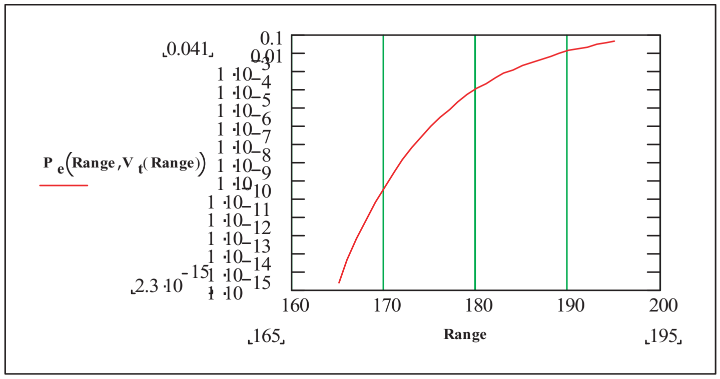

Decreasing the data rate from 1.25 Gbit/s to 100 Mbit/s increases the range for 10^{-12} BER by only about 30 meters in 200 dB/km attentuation

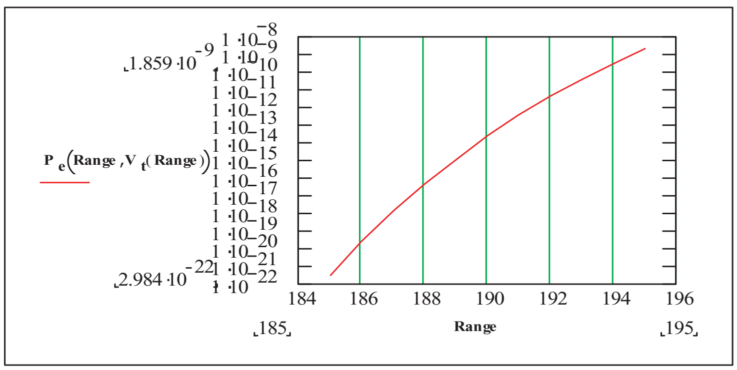

small prisms and lenses that deflect the light beam in to and out of the transmit path. The time scale of these fluctuations is about the time it takes a volume of air the size of the beam to move across the path and therefore is related to the wind speed. There are many theoretical papers written on this effect, primarily for ground to space applications, and, in general, are very complicated. It has been observed that for weak fluctuations, the distribution of received intensities is close to a log normal distribution. In the case of free space optics, which implies horizontal path propagation and therefore stronger scintillation, the distribution tends to be more exponential.

One parameter that is often used as a measure of the scintillation strength is the atmospheric structure parameter or C_n^2. This parameter roughly measures how turbulent the atmosphere is and is directly related to wind speed. *Figure 12* illustrates some measurements of this parameter taken at our facility in San Diego, California. The most interesting feature of the data is the fact that the parameter changes by more than an order of magnitude over the course of a day, being the worst or most scintillated during midday when the temperature is the greatest. The variation in C_n^2 is interesting because C_n^2 can be used to predict the variance in intensity fluctuations at the receiver through the formula displayed at the top of *Figure 13*. The variance is linearly proportional to C_n^2, nearly linearly proportional to 1/wavelength, and nearly proportional to the square of the link distance. Two interesting facts can be inferred from this. First, longer wavelength systems have proportionately larger variance in the scintillation intensity. For example, a system operating at 1550 nm has about twice the variance as a system operating at 780 nm. Second, the effect increases severely with range. At twice the range, the variance is four times greater. The graphs represent the variance as a function of range for the minimum, median, and maximum values of C_n^2 measured in *Figure 12*. The effects of scintillation are graphically illustrated in *Figure 14*. The figure depicts a receiver aperture with dark and light “”speckles” distributed randomly over the aperture. The size of these speckles scales like $(\mu R)^{1/2}$. Again, there are two interesting features to note about this scaling. First, the longer the wavelength, the larger the speckle. This is not good for system performance, because a smaller number of speckles at the receiver aperture means that less aperture averaging over the various intensities of each speckle can occur. If the receive aperture is only collecting one speckle, for example, then the required increase in system transmit power to make sure the BER is maintained even for a “dark” speckle is very large. Second, the size scales like the square root of the link distance; longer links imply larger speckles, which impacts system performance as previously described.

FIGURE 12

Scintillation and Measured Cn2 in San Diego, CA

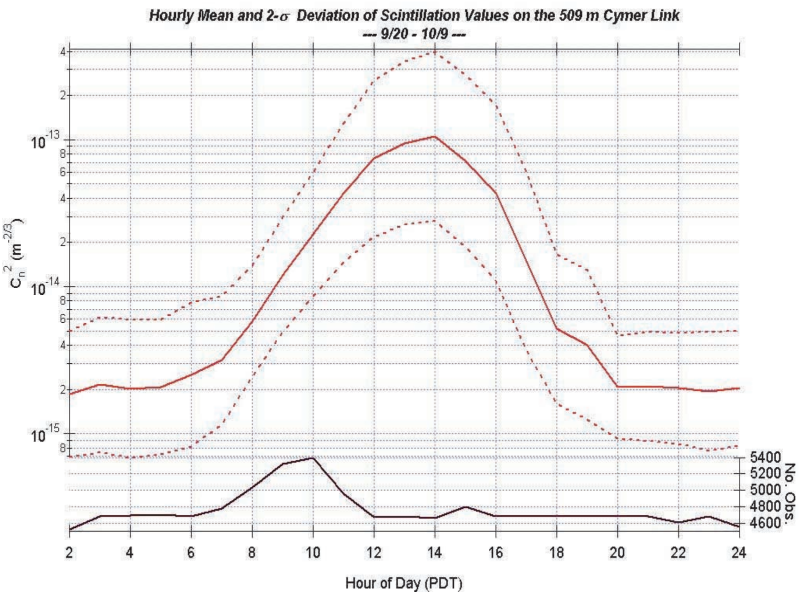

Figure 15 shows the effects of aperture averaging on the required link margin for a three-inch aperture. The plots are for a 100-meter link and a 1,000-meter link. The 100 meter link shows a received intensity distribution that is nearly Gaussian, centered about the mean transmit intensity of 1. Here, scintillation plays no role. In fact, for nearly *all* carrier-grade links (ranges of less than 500 meters), scintillation effects are a few dB at most. At 1 km, the intensity distribution is more skewed and requires about 13 dB more power for the same BER as the 100-meter link (to overcome scintillation, not geometrical effects). The point of this is that for *carrier-grade* links, scintillation is not an issue and therefore does not need to be addressed in the system design. For long links, there are system design choices that can be made—for exam-

FIGURE 13

Intensity Variance Derived from Cn2

$$\sigma_\chi^2\left(C_{nsq},k,L\right) := 0.31\cdot C_{nsq}\cdot k^{\frac{7}{6}}\cdot L^{\frac{11}{6}}$$

Variance in intensity fluctuation is derived from atmospheric structure constant

Scintillation effects scale (become worse) like Range^2

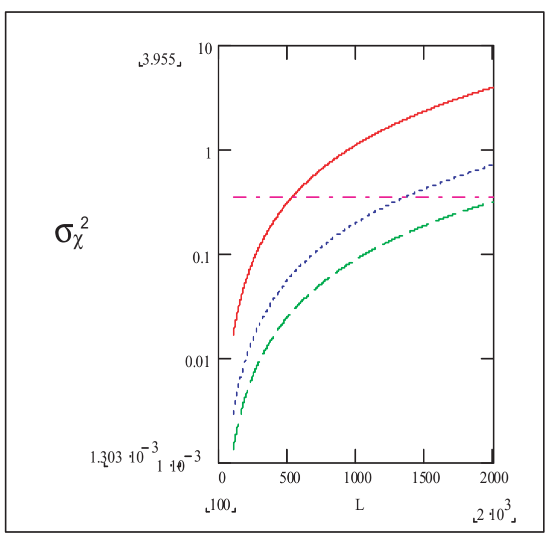

FIGURE 14

Scintillation Spots and Aperture Averaging

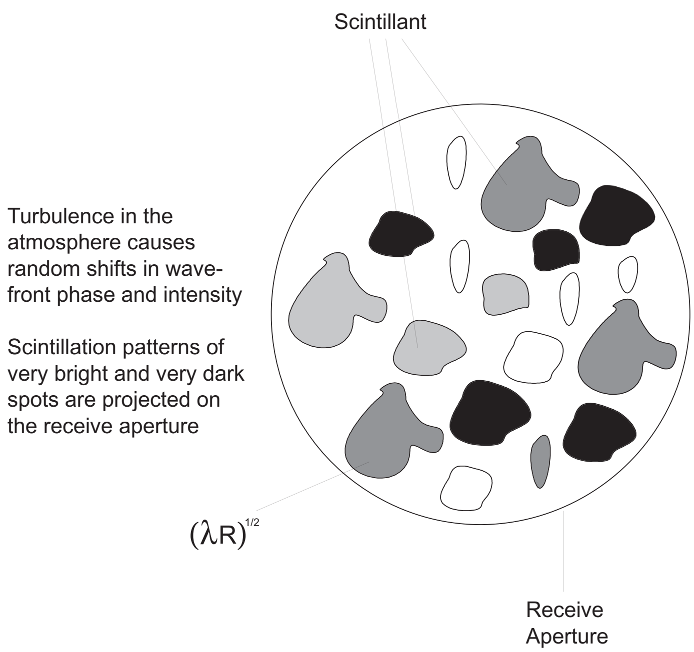

ple, using multiple laser transmitters—that substantially reduce scintillation effects, but again this is not necessary on a carrier-grade system, since the link ranges are generally less than 500 meters. Additionally, for carrier-grade systems, the link margins are engineered to take care of severe attenuation conditions, such as fog, and therefore already have enough margin to take care of any scintillation conditions, since the air is always clear when it is highly scintillated.

FIGURE 15

Theoretical Aperture Averaging Gain

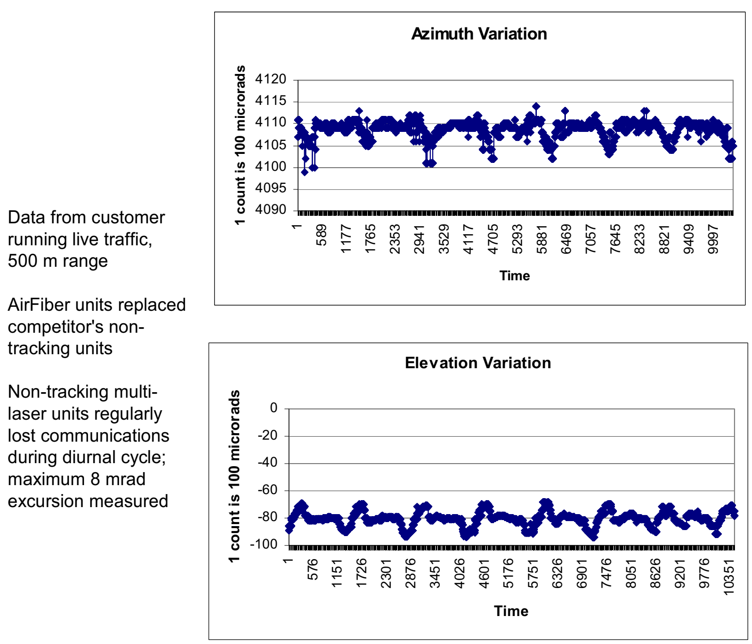

Laser reliability is an issue for carrier-grade free space optical systems that need to have an MTBF of eight years or more. There are basically two factors that influence the lifetime of a semiconductor diode laser: the average operating temperature of the case of the diode and the average operating output optical power of the diode. For AlGaAs diode lasers, the activation energy is about 0.65 eV, which implies that the lifetime of the diodes increases by about a factor of two for every 10 C decrease in the average operating temperature of the diode case. This is illustrated in the upper graph of *Figure 16*. Here, the lifetime increase as a function of case temperature is shown. For outdoor equipment, a general rule of thumb for most places in the world is that an average temperature of 25–30 C can be used to predict thermal effects on lifetime. More important than thermal effects is the average output power of the laser. In this case, the lifetime scales inversely with the cube of the power, as illustrated in the lower graph of *Figure 16*. This means that significant lifetime increases can be obtained in systems that employ automatic power control (by this we mean reduction in output laser power, not attenuation of power using a filter, for example) in their operation, since most of the time the air on a given link is clear and the output power can be reduced substantially. Systems that do not employ power control will probably have very difficult time meeting MTBF expectations of carriers.

Eye safety is also a concern among large service providers. There are basically two classes of eye safety to which free space vendors build: Class 1 and Class 1M. Basically, Class 1 means the beam can be viewed under any condition without causing damage to the eye. Class 1M means the beam can be viewed under any condition with the *exception* of an aided viewing device—for example, a binocular or other optical instrument. In terms of wavelength, the IEC60825 rulings permit 1550 nm devices to output about 50 times the flux level of near IR devices, such as 780 nm lasers. While this can be an advantage, the detectors at 1550 nm are noisier, and therefore, on a system level, there is a slight advantage for link margin, usually at a *substantial* increase in cost, especially for carrier-grade systems where the atmosphere is the limiting variable in the link equation, as previously discussed.

In summary, free space optical systems can be enhanced to carrier-grade (99.9 percent or better). This requires several simple design rules to follow for the vast majority of cities on the planet.

Notes

1. Atmospheric physics *fundamentally* limits range to less than 500 meters. *Accept* and *design* the hardware to meet this limit.
2. The equipment must be hardened to withstand the rigors of an outdoor environment; it *must* meet Telcordia GR-63, GR-487, and NEBS Level 3 product integrity standards.
3. Carrier-grade systems *must* have tracking. This is the *only* way to guarantee link margin.
4. Carrier-grade systems *must* have a carrier-class EMS.
5. Carrier-grade systems *must* meet Class 1M or Class 1 laser eye-safety standards as defined in IEC 60825 rulings.
6. Carrier-grade systems *should* have an optical management channel. This permits the carrier to avoid building an overlay network to manage the free space optical units.

FIGURE 16

Laser Lifetime versus Temperature and Power

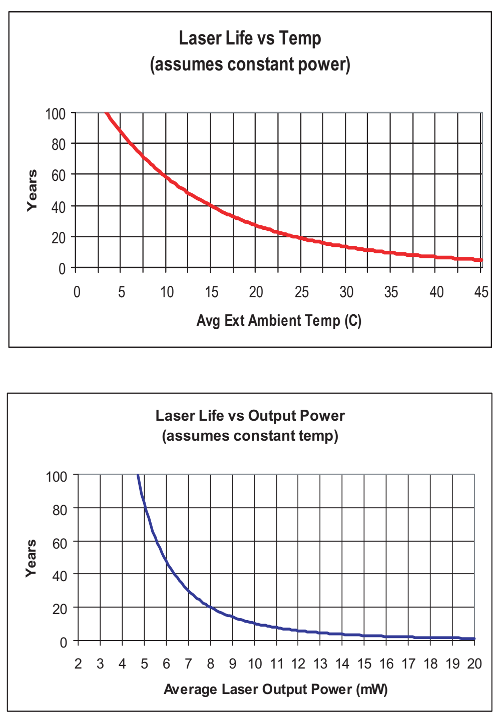

Laser lifetime approximately doubles for every 10°C decrease in case temperature (activation energy 0.65 eV)

Laser lifetime scales approximately as $1/P^3$. For example, a diode rated at 40 mW maximum output power gets a lifetime boost of 64x if run at 10 mW

Enhanced Services in the New Converged Telecom Environment

Peter Briscoe
President, Founder, and Chief Executive Officer
Convedia

Introduction

Telephony service is converging with Internet protocol (IP) and the Internet. This evolution will involve 15 years of transition in which services will undergo changes that will open up service design to a greater pool of creative talent. At present, the world is just starting on the long road to convergence (see *Figure 1*), and all players agree that it is not if but rather when this will happen.

In the past, telephony services were defined by committee, developed by large companies, deployed by big service providers, and hampered by a poor human interface.

The Internet is an open system, with end users and service providers able to create, share, and sell their products over the Internet as fast as technology development allows. The Internet has always been asynchronous with best-effort delivery.

The challenge looking forward is to ensure that the result of the convergence is the best of both worlds. The best of both worlds would see a reliable system with an open architecture that enables flexible service deployment options and regulations that encourage new business models for approaching telephony services and applications.

The intent is to unlock the telephony industry by empowering the end user (business or home) to rapidly create and deploy services—opening up telephony to the massive talent available in the world to find and develop new ideas for services and applications, resulting in increased wealth creation.

Enhanced Services

This discussion will focus on enhanced services and the issues involved with evolving to the new converged world. Enhanced services are defined as services that require technology assistance and typically use special software and/or hardware to interact with the users to provide the service. Enhanced services may be provided by terminal equipment, by enterprise-based equipment, or by network service provider (NSP)–based equipment (see *Figure 2*). Enhanced services are provided in the network when there are advantages due to centralization, such as processing power requirements, sharing of the resource to bring down cost, bandwidth restrictions, or cost of bandwidth. Enhanced services move to the edge when the technology is a commodity and the cost is minimal.

A good example of this is voice messaging, which started off as a terminal device application but is now available on many home phones, through the private branch exchange (PBX) or from the traditional telephony service provider. Conferencing is another feature that is available today through a residential service offered by the traditional service provider, through the PBX or through third-party conferencing service companies.

The present structure of enhanced services is shown in *Figure 3*. The terminal is simple with rudimentary signaling capability through dual-tone multifrequency (DTMF). The enhanced services platform consists of software service control logic and hardware processing of the audio streams.

Today, a separate hardware/software platform is used for each service and is called a feature server. Some common services that are provided by feature servers today are voice-mail, voice messaging, interactive voice response (IVR) services, conferencing services, unified communications services, and voice portal services

The evolution to a converged network and services, with the combination of the highly graphical interface of the Web and telephony capability, will result in new, exciting enhanced services and changes in how they are created, deployed, and used. The focus of this discussion is the enhanced service structure and how it changes as telephony and IP converge. This raises some fundamental questions:

- Which enhanced services are successful today and why?
- How should enhanced services be deployed in the converged network?
- Who will provide enhanced services in the converged environment?

FIGURE 1
Convergence

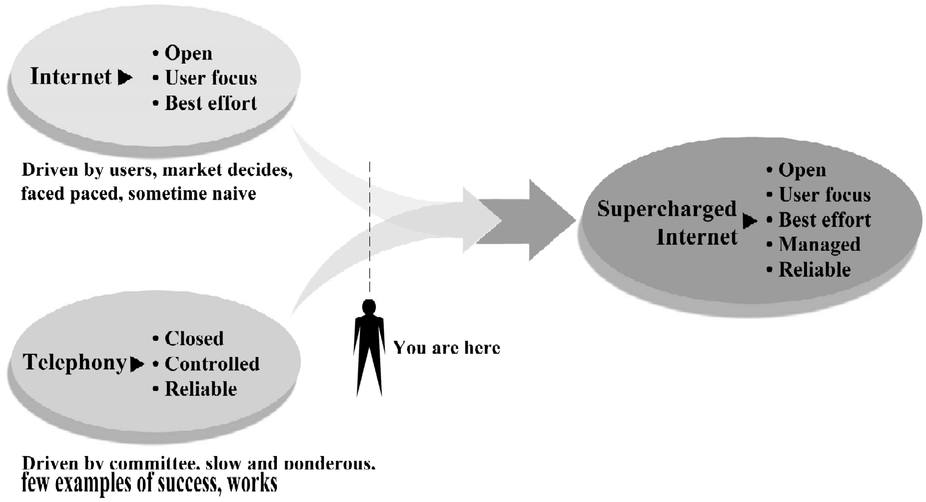

- Do service-creation environments make sense in a Web and rich media model?

This paper discusses these issues.

What Enhanced Services Are Successful Today and Why?

The following are the top successful enhanced services. Overall service revenues for these are in the billions-of-dollars range, with a 30 percent to 40 percent annual growth rate and projected revenue of tens of billions of dollars by 2005. Profitability is very high for these services.

- IVR systems (e.g., touch banking, touch reservations) provide convenience and save time by allowing people to access information and undertake activities from any phone without having to go to the institution where the service is provided. IVR systems also save businesses money in operator costs and lost revenue.
- Calling-card services save time and frustration by allowing access to phones without having cash on hand.
- Voice messaging and voice-mail saves time by allowing people to accomplish a task without having to be in touch with the other party. At this point in time, it is tightly coupled to the local access service provider.

FIGURE 2
Today's Enhanced Services

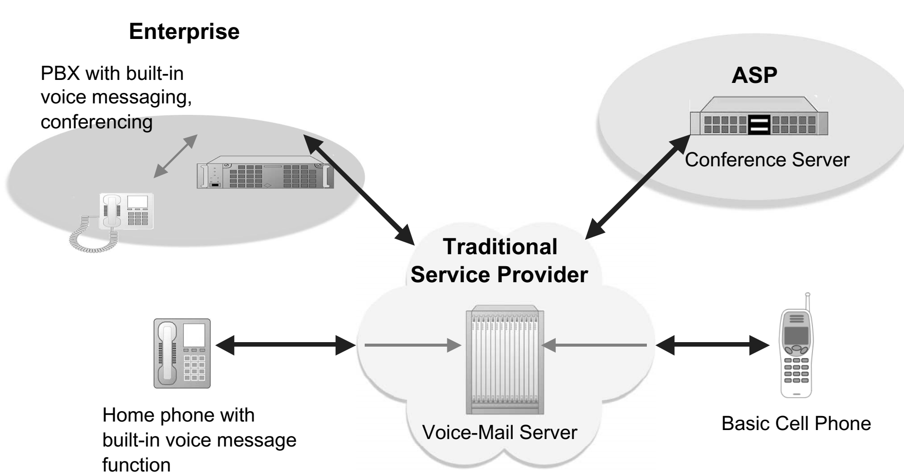

FIGURE 3

Enhanced Service Components

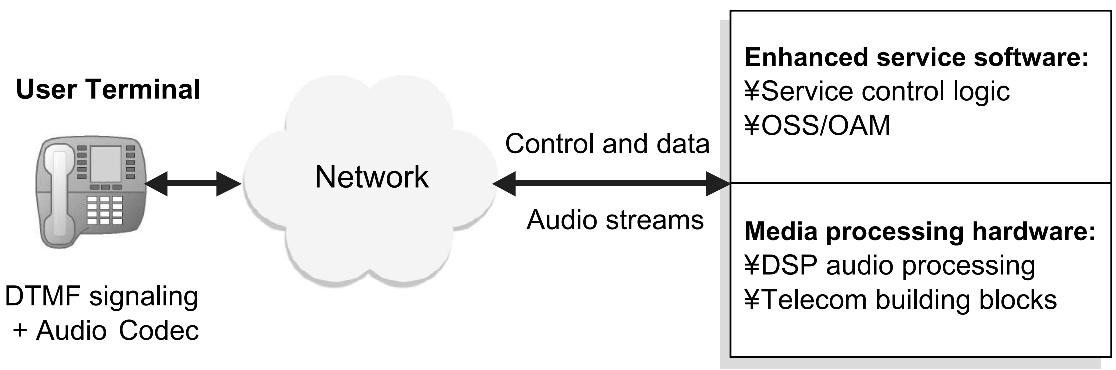

- Conferencing saves people time and money by allowing them to have meetings without traveling to a central location.
- Call centers give people access to assistance while lowering the cost to the business that must provide this.
- Videoconferencing is still in its infancy but has seen growth after the September 11th disaster.

The aforementioned services are the first targets of convergence due to their proven attraction and resulting revenue potential. The following services are new to the market and have not yet started to grow in any perceptible way in the marketplace.

Voice portals are slowing gaining popularity but have been trying to gain success for a few years now. This is a good example of how long it takes to bring a new service to market. Unified communications promises to be a valuable service, but it has not yet shown to be successful.

Enhanced services are extremely important to service providers. The movement to open competition and deregulation has put traditional service providers under a profit squeeze due to increased competition, causing basic voice service to become a commodity. This is shown by the two middle lines in *Figure 4*. These service providers must create a 12 percent return on invested capital to attract new capital infusions to continue evolution. With profit squeeze, this is becoming more and more difficult.

Service providers expect IP technology to provide as much as 40 percent savings over time division multiplexing (TDM) technology, as shown by the bottom curve in *Figure 4*. But it is enhanced services that hold the promise to large profits as shown in the upper curve. Some service providers already get up to 40 percent of their profits from enhanced service revenues. It is the enhanced services that are keeping service-provider profits at acceptable levels, and this is expected to become more significant in the future.

FIGURE 4

Profitability

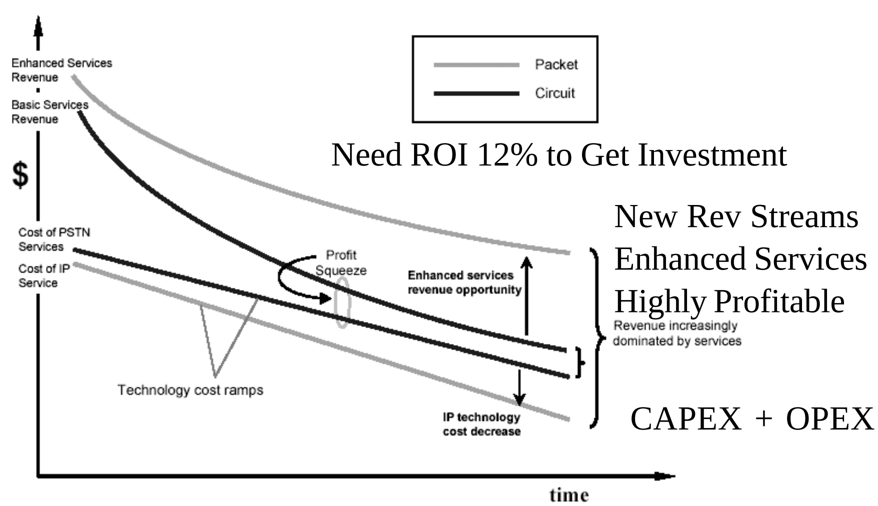

Source: Convedia, CIBC World Markets

The roll out of services will follow a fairly standards business model. The initial opportunities will be to replace the existing proven public switched telephone network (PSTN) services with IP–based services. This is the lowest risk and fastest return on investment (ROI), which is very important in the present investment climate. These services may get enhanced along the way with Web capabilities. As of yet, there has been no truly successful converged enhanced service where the power of the graphical Web interface is married to the quality and reliability of telephony.

How Should Enhanced Services Be Deployed?

It is nearly impossible to predict the next successful service. It is therefore important to create the right environment that will allow new services and applications to be explored so that the successful ones can become reality. The right environment includes an open architecture and access to new technologies, along with regulations that encourage people from industry and academia, as well as individuals, to apply their creative thinking and business acumen to new ideas for services and applications and then let the markets to dictate winners.

The softswitch architecture (SSA) provides a mechanism to ensure a seamless transition to the world of new media services (see *Figure 5*). The architecture has been formulated to take into account the business and technical realities of the world today and provide a flexible evolution path to get to the goal of a converged network and services.

There are many important attributes to this architecture, which provides the open environment required to stimulate new ideas, business models, and ultimately services. The SSA breaks down the telecoms service system into four distinct functional parts. These functions can be combined into products in many different ways that will allow different business requirements to be met and to explore entirely new business models.

One important change that the SSA has accomplished is to take the feature server of today and break it down into its software service logic and hardware media processing functional components called the application server and the media server. This step allows the software control logic for a service to be contained in a variety of devices leading to the very flexible deployment options that we seek for convergence. The complex, expensive hardware digital signal processor (DSP) equipment is reusable by multiple software services and applications among many users to bring the service CAPEX and OPEX down.

The DSP–audio processing building blocks located in the media server (bridging, tone detection and generation, announcement playing, recording and playback, speech recognition, and text to speech) are now basic telecom building blocks that can be accessed and controlled by any service software anywhere in the network through session initiation protocol (SIP) and voice extensible markup language (Voice XML), either directly or through a SIP proxy server (see *Figure 6*). This has the potential to unleash many new services if the service providers allow access to these blocks with suitable safeguards.

One obvious benefit is that the telecom building block control for services can be accessed through SIP, allowing telephony services to be more tightly integrated into computers and the Web and making telephony an integral part of computer applications for business and home.

There are a number of ways that enhanced services can be deployed:

1. Service platform
2. Enterprise server
3. End-user terminal

1. The service platform (see *Figure 7*) is a platform with all of the traditional telecom pieces: operations sup-

FIGURE 5

Softswitch Architecture

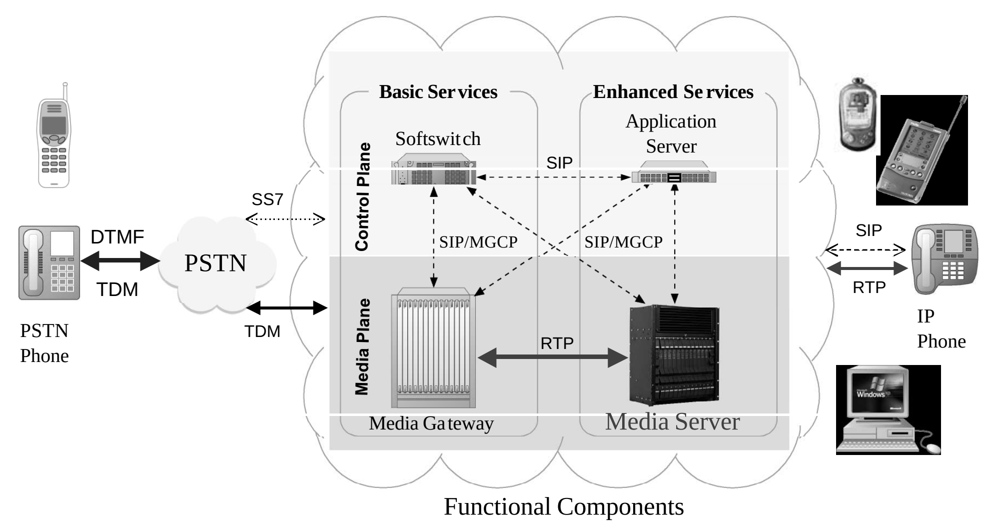

FIGURE 6

Media Server Building Blocks

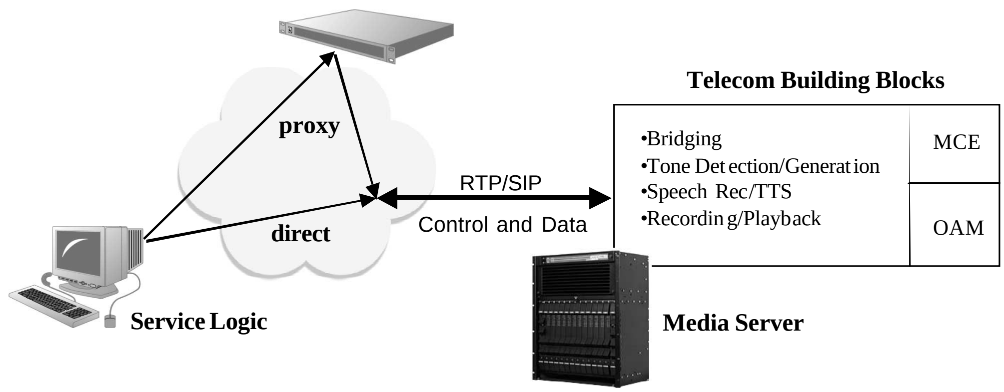

port systems (OSSs), billing, policy, application plug-ins, media server control, and a service-creation environment (SCE). This is an expensive product but very flexible and complete.

Service platforms are offered in various levels of complexity and sophistication. Some include all of the functions shown in *Figure 7* and allow other vendors to plug their service modules into the platform. Some service platforms have imbedded predeveloped services and do not contain the OSS/business support system (BSS) capabilities in-board but pass information to other existing systems. These tend to be much cheaper than the full platforms but lack the sophistication. .

2. Enterprise PBXs can have service control built into their applications in order to permit the automated control of services by the enterprise equipment to provide services to the enterprise users in conjunction with the service provider's services and network. In this scenario, the enterprise device contains the application server functionality and controls the network service on behalf of its enterprise terminals. This may be used to provide new services to old PBX terminals or to translate PBX feature-rich services into simpler network-based services.

   This approach allows the enterprise to retain control of its service design and still use the services and reach of the large service provider.

3. The Internet model has different characteristics than the TDM model. Peer-to-peer capability in IP allows basic calling without a "call control" being required (see *Figure 9*). The only issue in crossing network boundaries is to maintain network quality of service

FIGURE 7

Services Platform

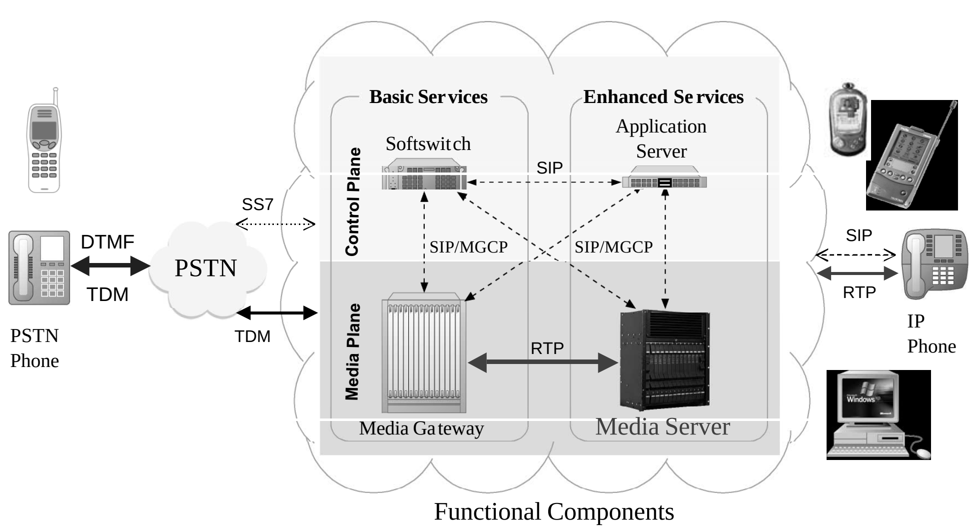

FIGURE 8
Enterprise Server

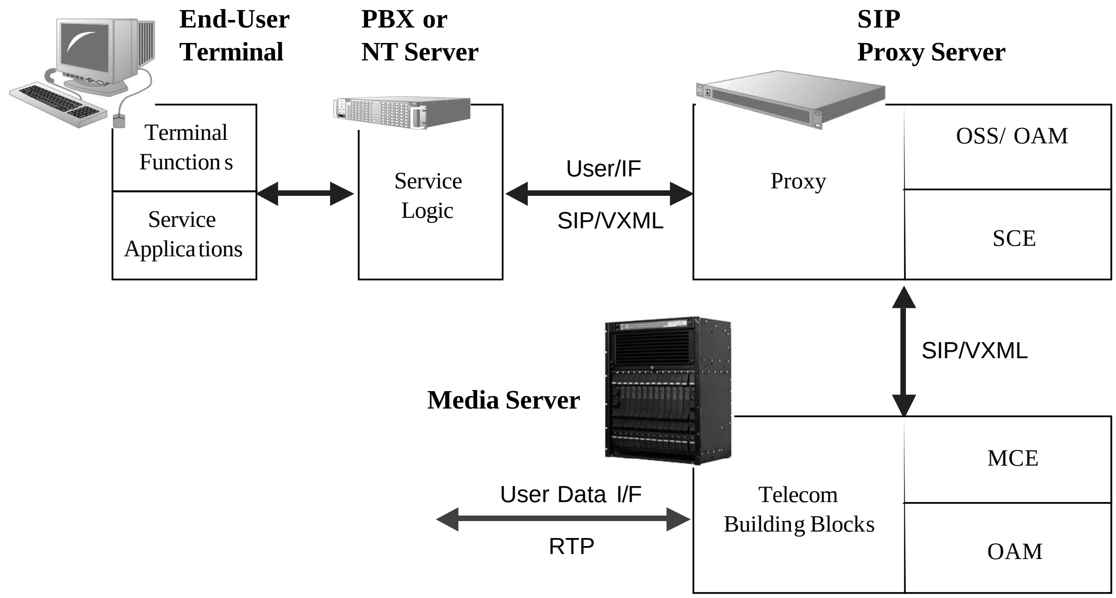

(QoS) throughout the network. The opening up of the service control interface also allows service providers to offer the telecom service building blocks to end users to create their own services or imbed services into their own computer applications.

SIP, along with service control languages such as Voice XML, allow a user to do more than request services; the user can activate and control services from the terminal equipment. End-user terminal devices, such as personal digital assistants (PDAs) and SIP phones, can have service applications built into them and allow the end user to program and control their own services. This puts service creation and control into the hands of the end users and application developers, similar to what happened with the PC and with the Internet. When the creativity of tens of thousands of programmers is let loose on telephony services and the promise of great wealth to the successful ones, there will be many exciting new services created and huge traffic generated.

FIGURE 9
End-User Terminal

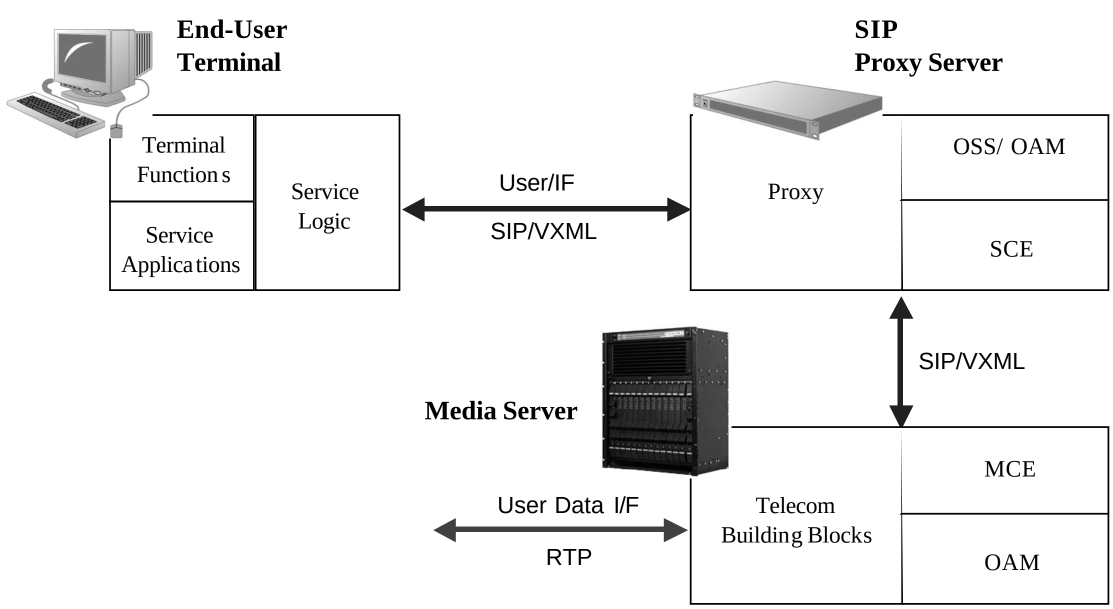

Undoubtedly, systems 2 and 3 are a clear departure from the normal service environment of the present telephony system and will require a number of changes. Security and authentication is critical before giving an end user access to the service building blocks. Billing for this style of service can be monthly or by the session and service type. Predefined and tested Voice XML scripts will have to be published to users to enable specific services until the development community obtains widespread knowledge. But, of all of the methods, this unleashes the most creativity, albeit with the most change as well.

Who Will Provide the Services?

Working from the core of the network outward, we see that there are multiple players that have built viable businesses at natural interface points of the network. From experience, it is likely that most, if not all, of the existing business models will continue to grow and flourish even as new players enter the market. Service providers all over the world understand the strategic importance and cost benefits of moving forward quickly with convergence.

In this section, we explore the deployment of voice over IP (VoIP) by various services providers, shown in *Figure 10*, including local providers (ILEC, wireless, and cable multiple-system operators [MSOs]), long-distance providers (IXCs and next gens), and enterprises.

Long-Distance Service Providers

Long-distance (LD) service providers have already started to adopt VoIP for LD arbitrage and LD Class-4 replacement. Many of these providers have built massive IP backbone networks, making them ideally positioned to provide VoIP services. These service providers wish to take advantage of their central position in the network to offer services that span geographic areas. They have the benefit of shorter backhauling of traffic to a central point and access to a high percentage of cross-country traffic.

Local Service Providers

There is a battle underway for ownership of the access for home and business services. Presently, the ILEC or the Post Telephone and Telegraph Administration (PTT), have control of telephony services. But challenges to this dominance can occur from the cable-television (CATV) players and the wireless players. Traditional local service providers wish to add more services and increase their revenues and profits. In some service providers, enhanced services, such as voice-mail and card services, provide as much as 50 percent of the company profits. These players can not cannibalize their existing services or disrupt their amortization scheduled for their installed base of equipment, but they can use the converged services to attack market areas where they are weak or where they have little sunk investment. Rather than wait for this while VoIP starts without them, they are busy exploring areas that are good candidates for new revenue streams.

IP Centrex

Enterprise organizations are outsourcing everything that is not core to their business. This provides an opportunity for incumbent players to gain back enterprise customers that they lost to PBX vendors in the 1980s. Hosted business services such as IP Centrex give the incumbent service providers a powerful story that is separate from their core residential

FIGURE 10

Service Options

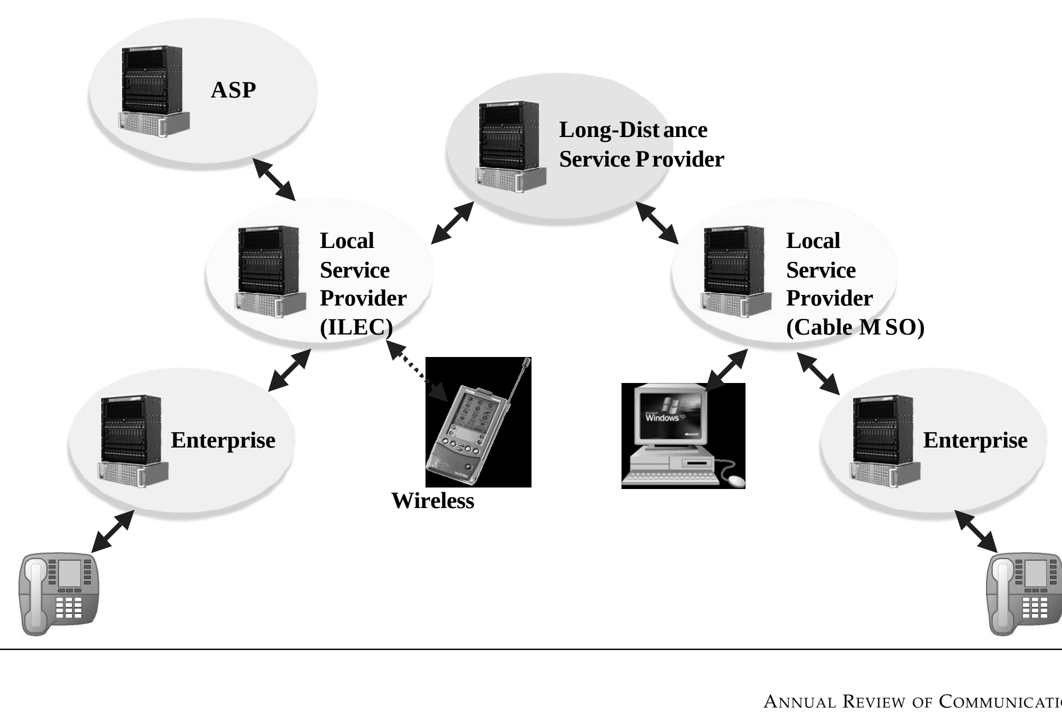

telephony business, allowing them to deploy VoIP, learn the technology, and position for convergence.

IP VPNs

Large players have deployed IP VPNs to service their largest customers. Today, those mainly carry data traffic, although many of them contain VoIP traffic hidden in the data. The service provider is ideally positioned to put an enhanced services platform on this core network and to add telephony services on these IP VPNs. This protects the enterprise customer from newcomers with VoIP, and it does not require a complete change over of the service provider's residential TDM equipment.

Enhanced Services Offload

ILECs can use enhanced service offload to consolidate enhanced services onto a single technology. This technique requires a mini softswitch package of products that connect into the TDM network at a logical trunking point. This package consolidates all of the traditional TDM feature server products into a single flexible, scalable product that is IP–enabled. This provides cost savings in both CAPEX and OPEX over TDM–based solutions while capping purchases of dead-end TDM equipment and positioning for VoIP. The service provider can even offer Web-based enhanced services off the same platform to explore these new, exciting services.

Cable MSOs

The CATV industry is poised to provide new, exciting services, being the only other player with substantial bandwidth capability to the home. The new entrants into local access are looking to replicate the Class-5 features at a lower cost than the incumbents in order to offering new lower-cost service that will help them to capture customers away from the incumbent players. Wireless players are hungry for new, exciting services to differentiate themselves, and with a wave of intelligent wireless handsets on the way, they will have the graphical user interface (GUI) to do this.

Application Service Providers (ASPs)

ASPs are relatively new service-provider business models. These players rent a connection to the service provider's network and sell services across the network to the consumers or small business. In some cases, they supply the local or LD service provider with services, and sometimes they compete with the local and LD service providers and rely on their ability to provide richer service features, faster time to market, and better customer support than the larger service providers.

Enterprise

Enterprises may wish to control and manage their services locally. This can be done by a PBX or by a PC running Windows XP. They can look inward and outward to provide enterprise-wide services that are hosted by the enterprise server or through a combination of enterprise services and carrier-hosted services.

End Users

Consumers and small businesses may wish to control basic services directly from their PCs or intelligent terminal devices. They may imbed the application server functionality in the intelligent terminal device to allow service control and management to be built into the terminal application. In this scenario, the terminal device may access complex media server processing functionality in the network and pay on a monthly basis or as required.

Summary

Telephony service will be converging with the Internet during the next 15 years. Enhanced services are one of the keys to profitable business in telephony. Enhanced services can be provided through service platforms to end users. Service providers are already starting to deploy enhanced services where economics show there is a benefit, such as for Class 4, IP Centrex, enhanced services offload, and IP VPNs. In the near future, enhanced services building blocks can be provided by the network for any user to use through SIP and Voice XML.

The Battery and the Digital Load

Isidor Buchmann
Founder and Chief Executive Officer
Cadex Electronics, Inc.

With the move from analog to digital devices, new demands are being placed on the battery. Unlike analog equipment that draws a predictable and steady current, digital devices load the battery with short, high current bursts.

One of the urgent requirements of a battery for digital applications is low internal resistance. Measured in milliohms (mΩ), the internal resistance is the gatekeeper that, to a large extent, determines the runtime. The lower the resistance, the less restriction the battery encounters in delivering the needed power bursts. A high mW reading can trigger an early "low battery" indication on a seemingly good battery because the available energy cannot be fully delivered.

In this article, we examine the current requirements of analog and digital communications devices. *Figure 1* provides typical examples of peak current of the analog two-way and digital Tetra radio, as well as the AMP, Global Systems for Mobile Communications (GSM), time division multiple access (TDMA), and code division multiple access (CDMA) mobile phones.

Why Do Seemingly Good Batteries Fail on Digital Equipment?

Service technicians have been puzzled by the seemingly unpredictable battery behavior when powering digital equipment. With the switch from analog to digital wireless communications devices, particularly mobile communications equipment, a battery that performs well on an analog system may show irrational behavior when used on a digital unit. Testing these batteries with a battery analyzer produces good capacity readings. Why then do some batteries fail on digital devices but not on analog?

The overall energy requirement of a digital mobile phone is less than that of the analog equivalent, however, the battery must be capable of delivering high current pulses that are often several times that of the battery's rating. Let's look at the battery rating as expressed in C rates.

A 1C discharge of a battery rated at 500 mAh is 500 mA. In comparison, a 2C discharge of the same battery is 1,000 mA. A GSM phone powered by a 500 mA battery that draws 1.5A pulses loads the battery with a whopping 3C discharge.

A 3C rate discharge is acceptable for a battery with very low internal resistance. However, aging batteries, especially Li-ion and NiMH chemistries, pose a challenge because the mΩ readings increase with use. Improved performance can be achieved by using a larger battery, also known as an extended pack. Somewhat bulkier and heavier, an extended pack offers a typical rating of about 1,000 mAh or roughly double that of the slim-line. In terms of C rate, the 3C discharge is reduced to 1.5C when using a 1,000 mAh instead of a 500 mAh battery.

As part of ongoing research to find the best battery system for wireless devices, Cadex has performed life-cycle tests on various battery systems. In *Figure 2* through *Figure 4*, we examine NiCd, NiMH, and Li-ion batteries, each of which generate a good capacity reading when tested with a battery analyzer but produce stunning differences on a pulsed discharge of 1C, 2C, and 3C. These pulses simulate a GSM phone.

A closer look reveals vast discrepancies in the mΩ measurements of the test batteries. In fact, these readings are typical of batteries that have been in use for a while. The NiCd shows 155 mΩ, the NiMH shows 778 mΩ, and the Li-ion shows 320 mΩ, although the capacities checked in at 113 percent, 107 percent, and 94 percent, respectively, when tested with the DC load of a battery analyzer. It should be noted that the internal resistance of a new battery reads between 75 mΩ and 150 mΩ.

From these charts, we observe that the talk time is in close relationship with the battery's internal resistance. The NiCd produces a long talk time at all C rates. In comparison, the NiMH only works at a lower C rate. The Li-ion performs better but is marginally at a 3C discharge.

How Is the Internal Battery Resistance Measured?

A number of techniques are available to measure the internal battery resistance. One common method is the direct current (DC) load test, which applies a discharge current to the battery while measuring the voltage drop. Voltage over current provides the internal resistance.

The alternating current (AC) method, also known as the conductivity test, measures the electrochemical characteristics of a battery. This technique applies either a fixed frequency or a frequency range from 10 Hz to 1,000 Hz to the battery terminals. The impedance level affects the phase shift between voltage and current, which reveals the condition of the battery.

FIGURE 1

Peak Power Requirements of Popular Global Mobile Phone Systems

	Two-Way Radio	Tetra Radio	AMP Mobile Phone	GSM Mobile Phone	TDMA Mobile Phone	CDMA Mobile Phone
Type	Analog	Digital	Analog	Digital	Digital	Digital
Peak Power	2–4 Watts	1 or 3 Watts	0.6 Watts	1–2 Watts	0.6–1 Watts	0.2 Watts
Peak current	1.5A at 4W	1 or 3A	0.3A DC	1–2.5A	0.8–1.5A	0.7A
In service since	1962	1997	1985	1986	1992	1995

Moving from analog to digital equipment reduces the overall energy need but increases the peak current during load pulses. The wattage varies with signal strength.

Some AC resistance meters evaluate only the load factor and disregard the phase-shift information.

Cadex uses the discreet DC method to measure internal battery resistance. Added to the Cadex 7000 Series battery analyzers, a number of charge and discharge pulses are applied, which are scaled to the mAh rating of the battery tested. Based on the voltage deflections, the battery's internal resistance is calculated. Known as *Ohmtest*™, the mΩ reading is obtained in five seconds.

None of the three methods is dead accurate. The discrepancies are reasonably small on a good battery, but the readings get more diverse on weaker packs. *Figure 5* compares the accuracy obtained using the three methods.

Resistance measurements alone do not provide a reliable indication on the battery's performance. The mΩ readings may vary widely depending on battery chemistry, cell size (mAh rating), type of cell, number of cells connected in series, wiring, and contact type.

When using the impedance method, a battery with a known performance should be measured and its readings used as a reference. For best results, a reference reading should be on hand for each battery type. *Figure 6* provides a guideline for digital mobile phone batteries based on impedance readings.

The milliohm readings are related to the battery voltage. Higher voltage batteries allow higher internal resistance before the system fails, because less current is required to

FIGURE 2

Talk Time of a NiCd Battery under the GSM Load Schedule

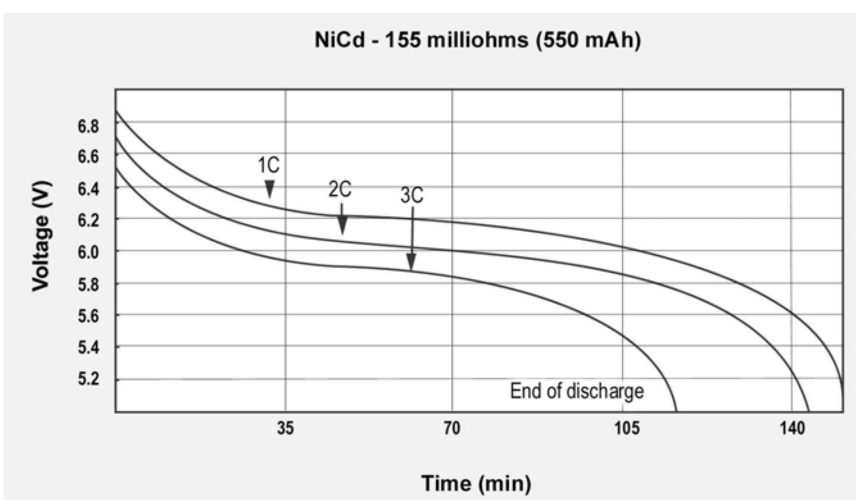

This battery has 113 percent capacity and 155mΩ internal resistance.

FIGURE 3

Talk Time of a NiMH Battery under the GSM Load Schedule

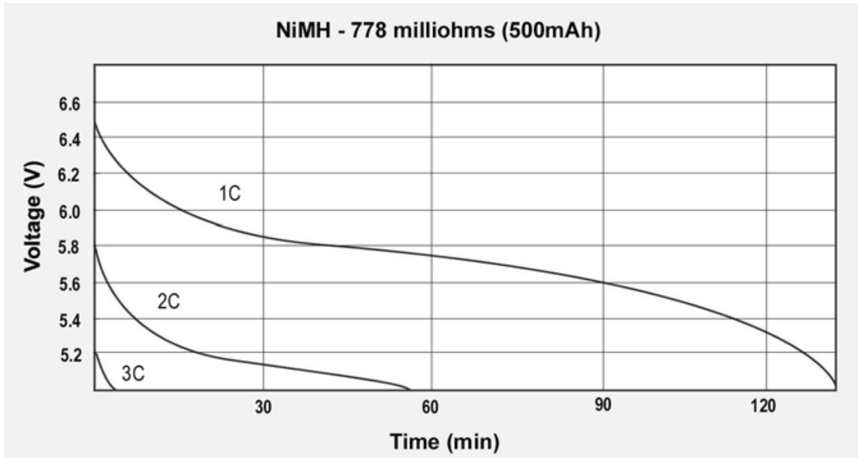

This battery has 107 percent capacity and 778mΩ internal resistance.

deliver the same power. The ratio between voltage and milliohm is not totally linear. There are certain housekeeping components that are always present whether the battery has one or several cells. These are wiring, contacts, and protection circuits.

Temperature also affects the internal resistance of a battery. The internal resistance of a naked Li-ion cell measures 50 mΩ at 25° C (77° F). If the temperature increases, the internal resistance decreases. At 40° C (104° F), the internal resistance drops to about 43 mΩ and at 60° C (140° F) to 40 mΩ. While the battery performs better when exposed to heat, prolonged exposure to elevated temperatures is harmful. Most batteries deliver a momentary performance boost when heated.

Cold temperatures have a drastic effect on all batteries. At 0º C (32º F), the internal resistance of the same Li-ion cell drops to 70 mΩ. The resistance increases to 80 mW at –10º C (14º F) and 100 mΩ at –20º C (–4º F).

The internal resistance readings work best with Li-ion batteries because the degradation follows a linear pattern with cell oxidation. The performance of NiMH batteries can also be measured with the internal resistance method, but the readings are less dependable. There are instances when a poorly performing NiMH battery can also exhibit a low mΩ reading.

Testing a NiCd on resistance alone is unpredictable. A low resistance reading does not automatically constitute a

FIGURE 4

Talk Time of a Li-ion Battery under the GSM Load Schedule

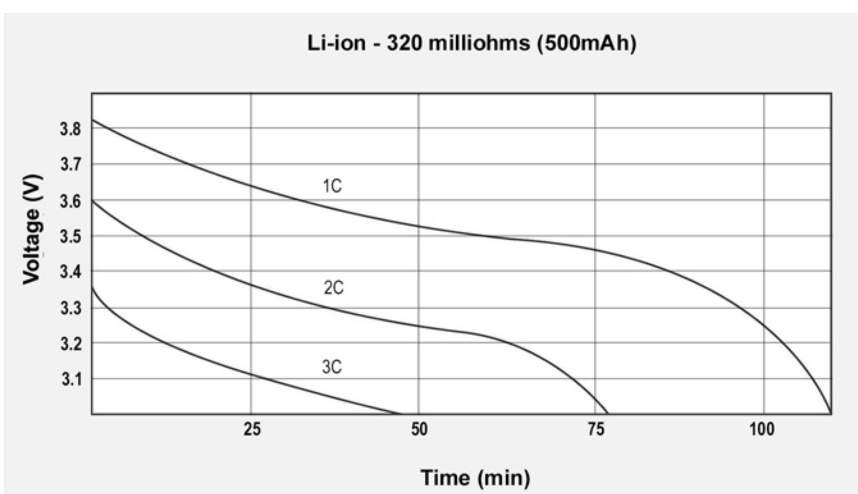

This battery has 94 percent capacity and 320mΩ internal resistance.

FIGURE 5

Comparison of the AC, DC, and Cadex Ohmtest™ Methods

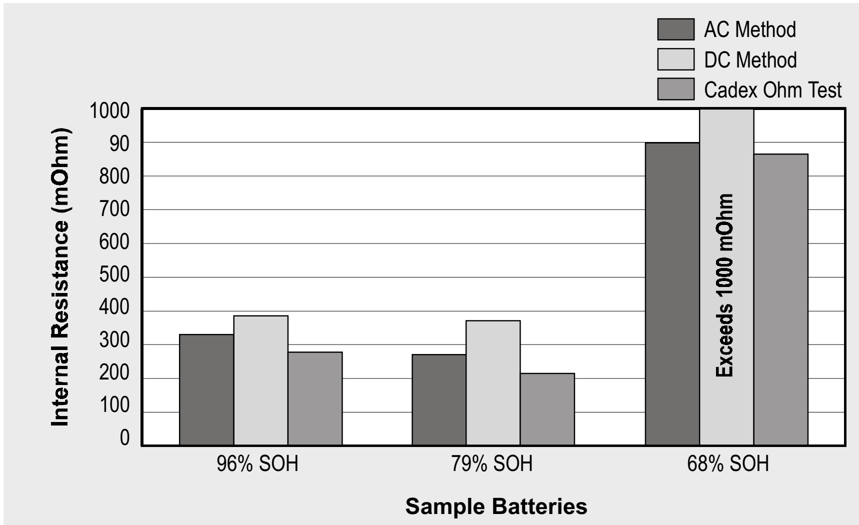

State-of-health readings were obtained using the Cadex 7000 Series battery analyzer by applying a full charge/discharge/charge cycle.

good battery. Elevated impedance readings are often caused by memory, a phenomenon that is reversible. Of course, high internal resistance can have sources other than memory alone.

Summary

Customer demand has compelled manufacturers to equip portable devices with batteries that provide a long talk time, are small, and are light in weight. By packing more energy into a pack, other qualities may be neglected, one of which is internal resistance and longevity.

Predictable low mΩ reading sand long service life is found in the NiCd family. This chemistry has been replaced with higher-energy dense batteries for many wireless applications. In addition, negative publicity about the memory phenomenon and concerns of toxic metals have caused a shift toward alternative choices.

For many applications—including biomedical devices, power tools, and most notably the Tetra system—the NiCd may be the only battery that has the endurance of delivering high pulse current under continuous usage. Other chemistries are simply too fragile. The resistance on a NiMH rises after a few hundred charge/discharge cycles. In comparison, a properly maintained NiCd provides over 1,000 cycles.

For many portable devices, the battery is one of the most expensive components. It is also the only part that repeatedly fails during the life of a product. It is therefore prudent that manufacturers not only focus on high energy density, but also address the issue of longevity. The longer a battery lasts, the fewer batteries are discarded—a win-win situation for business and the environment.

This article contains excerpts from the second edition of Batteries in a Portable World – A Handbook on Rechargeable Batteries for Non-Engineers.

FIGURE 6

State-of-Health on Mobile Phone Batteries Based on Internal Resistance

Milliohms	Battery Voltage	Ranking
75–150 mΩ	3.6V	Excellent
150–250 mΩ	3.6V	Good
250–350 mΩ	3.6V	Marginal
350–500 mΩ	3.6V	Poor
Above 500 mΩ	3.6V	Fail

The milliohm readings relate to the battery voltage; higher voltage allows higher milliohm readings.

Data on MMS: An Introduction to Mobile Multimedia Messaging

Simon Buckingham
Chief Executive Officer
Mobile Streams

1. Evolution from Text to Multimedia

Over time, the nature and form of mobile communication is getting less textual and more visual. Mobile messaging is evolving beyond text by taking a development path from short message service (SMS) to enhanced messaging service (EMS) to multimedia messaging service (MMS). Mobile Streams already publishes its renowned SMS reports called "SMS Express" and "SMS Tech," and its EMS report "Yes 2 Enhanced Messaging."

The transition that we will see from SMS to MMS is akin to the revolution from DOS to Windows in the computing world. It was this change that took computing from the early adopter innovator category into the early majority and onwards into the late majority mainstream status. When you compare a text message or an operator logo with a multimedia message or screensaver, the difference is an order of magnitude. The main features of this transformation from text to multimedia are shown in *Table 1*.

The SMS is the ability to send and receive text messages to and from mobile telephones. The text can comprise of words or numbers or an alphanumeric combination. SMS was created when it was incorporated into the Global System for Mobile Communications (GSM) digital mobile phone standard. The first short message was sent in December 1992 from a personal computer (PC) to a mobile phone on the Vodafone GSM network in the United Kingdom. Two-way SMS is supported on GSM, code division multiple access (CMDA), and time division multiple access (TDMA) networks.

The EMS is the ability to send ringtones, logos, and pictures to mobile phones and also to send and receive a combination of simple media such as melodies, pictures, sounds, animations, modified text, and standard text as an integrated message for display on an EMS compliant handset. There are many different potential combinations of these media. For example, when an exclamation mark appears in the enhanced message, a melody could be played. A simple black and white image could be displayed along with some text and this sound effect. EMS is an enhancement to SMS but is very similar to SMS in terms of using the store-and-forward SMS centers, the signaling channel, and the like to realize EMS. There are no network modifications needed to support EMS. Only new phones supporting EMS are needed. The first EMS compliant handsets became available in mid 2001. EMS is a GSM and CDMA standard.

The MMS is, as its name suggests, the ability to send and receive rich media messages comprising a combination of text, sounds, images, and video to and from MMS–capable handsets. The MMS confers the ability to send still images such as mobile postcards, mobile pictures, mobile screensavers, mobile greeting cards, mobile maps, and business cards. Additionally, moving images, cartoons, and interactive video will also be supported by multimedia messaging. New mobile network infrastructure is needed for multimedia messaging—in addition to implementing the new bearer services such as general packet radio service (GPRS) and third generation (3G), new network elements such as multimedia messaging relays and stores will be needed. MMS is a service that will run over Internet protocol (IP)–based mobile networks—initially GPRS networks—and will later be ported to enhanced data rates for GSM evolution (EDGE) and 3G networks.

2. Comparison between SMS and MMS

The SMS is a very popular mobile messaging service, with tens of billions of SMS messages sent and received each month globally. MMS will want to tap into and replace some of the SMS traffic over time. It is therefore useful to compare SMS and MMS. SMS and MMS share some similarities and have some discontinuities, as detailed in *Table 2*.

Both SMS and MMS are non–real-time services—this means that there is an intermediate platform such as the SMS center (SMSC) or the MMS relay through which the short or multimedia messages pass. Another characteristic that SMS and MMS hold in common is that both include confirmation of message delivery—the sender of the message can find out whether or not the message they sent was successfully delivered.

Media Supported
The SMS supports text and binary as media, allowing, for example, rudimentary images to be sent and received. The

TABLE 1

Type	Characteristics	Content Reformatting for Mobile Necessary?	Applications	Support	Time Frame for Availability
SMS/Text Messaging	100–200 characters	Yes	Simple person-to-person messaging	All phones	1990s
Nokia Smart Messaging	Simple rudimentary images	Yes	Simple person-to-person messaging with a visual feel	Some networks and Nokia phones only	1999 onwards
Enhanced Messaging (EMS)	Text messages plus simple media formats, e.g., sound, animation, picture, text formatting enhancements	Yes	Simple person-to-person messaging with a visual feel	EMS standards expected to be widely adopted	2001 onwards
Multimedia Messaging (MMS)	Messages in multiple rich media formats, e.g., video, audio plus text	Sometimes	Sharing still images between phones and PCs and later simple person to person messaging with a visual feel	MMS standards expected to be widely adopted	2002 onwards
Source: Mobile Streams					

overwhelming majority of all SMS messages are pure and plain text, however. Multimedia messages, on the other hand, can be coded in various media from text to images to sounds to video clips to a combination of these. As such, the MMS is a much more complicated and powerful service that supports far more media and rich media. For this reason, SMS is to mobile phones what DOS was to PCs, whereas MMS is to mobile phones what Windows was to the PC. This is a revolutionary step, requiring EMS in the middle to steer an evolutionary migration path in mobile messaging. It also mandates terminal negotiation between network and phone to assess capabilities. SMS is the lowest common denominator that all phones support; MMS is a complex service where different phones will have varying media support. Whereas SMS is integrated into every mobile phone (whether that phone also supports EMS, wireless application protocol [WAP], or MMS), MMS clearly requires new terminals.

Delivery Mechanism

All short messages are sent and received over the signaling channel, a channel that is an additional transport mechanism on GSM networks over and above the radio channels themselves. The signaling channel is a little like the hard

TABLE 2

Feature	SMS	MMS
Store and Forward (Non–Real-Time)	Yes	Yes
Confirmation of Message Delivery	Yes	Yes
Communications Type	Person to Person	Application to Person Person to Person
Media Supported	Text plus Binary	Multiple: Text, images, video
Delivery Mechanism	Signalling Channel	Data Traffic Channel
Protocols	SMS Specific, e.g., SMPP	WAP and General Internet, e.g., MIME, HTTP, SMTP
Configuration	Simple Telephone Number	Diverse Parameters
Platforms	SMS Center	MMS Relay plus Others
Principle Applications	Simple Person to Person	Still images, person to person, server-based MMS services, e.g. video news
User Behavior	Discreet	Indiscrete
Source: Mobile Streams		

shoulder or on-ramp on a motorway/expressway—it runs parallel to the traffic lanes themselves. SMS can be transmitted concurrently to other data types—text messages can be sent and received while the user is also on a voice, circuit-switched data or fax call.

Multimedia messages, on the other hand, will be transmitted over the traffic channels themselves, where other data types from voice to data will also be transported. The high capacity of 3G networks will mean that all of these different traffic types can share the same radio resource without the likelihood of congestion. Using the traffic channel helps to overcome the capacity limitations of the signaling system 7 (SS7) signaling channel (see www.mobileSS7.net for more information).

In the GPRS world, therefore, multimedia messages will share limited network capacity with voice and other calls. This will have an important implication for quality of service (QoS), since multimedia-message delivery reliability will be affected if networks are congested for any and all types of mobile services.

Protocols

When SMS was standardized in the early to mid-1990s, the Internet was an obscure academic communications medium. The original European Telecommunications Standards Instititute (ETSI) specifications for SMS closely mandated some areas of SMS and left others open to competition. As a result, proprietary protocols were developed—every SMS–center vendor developed its own interface such that application developers needed to implement different interfaces when porting their applications and services to network operators that had different SMSCs. Short message peer-to-peer protocol (SMPP) has recently become the de facto SMSC interface protocol and is also likely to be used for certain MMS interfaces. Furthermore, an outdated protocol—X.25—remains a popular access mechanism for connecting applications to SMSCs.

MMS, on the other hand, came of age in the Internet world where open systems and standard protocols reign and a wide range of these protocols exist. There is therefore no need to reinvent the wheel—existing standard protocols can be used—and MMS needs to tap into the vibrancy and innovativeness of the Internet companies if it is to maximize its full potential. MMS uses standard Internet protocols such as multipurpose Internet mail extensions (MIME) and simple mail transfer protocol (SMTP) for access to the multimedia messaging service environment (MMSE). MMS is basically a presentation layer for basic e-mail protocols.

While the first MMS implementations will be based on enhanced WAP protocols (WAP MMS encapsulation), later MMS versions will also support non–WAP, standard Internet protocols for communication between terminal and MMS relay, such as hypertext transfer protocol (HTTP) over transmission control protocol (TCP)/IP.

Configuration

SMS is a completely simple service to use. There are no special numbers to remember—the same number that you use to call is the one you use to text, and no additional parameters need to be set by the end user to successfully send the text message. MMS, however, is a very complex service with device negotiation and capability recognition, and different classes of devices with different media support. In the initial implementation of MMS using WAP, WAP Push is used such that concatenated SMS messages are used to transport the notification data (sender, size, retrieval universal resource locator [URL], etc.) encapsulated in a WAP Push data unit. The third-generation partnership project (3GPP) IP–based implementation proposal for MMS does not include SMS notification—it assumes a pure HTTP payload between the terminal and the relay. The 3GPP has defined a frame that allows for different implementations but assumes that the IP address used is fixed. This means that a lot of IP addresses will be needed for MMS, which will be possible using only IPv6 and not the current IPv4.

This difference in the complexity of notifications between MMS and SMS is yet another crucial one that does not favor MMS.

Platforms

In SMS, the SMSC is the heart of the service, with all short messages of any type passing through an SMSC to and from mobile phones. As such, there is one platform type that dominates SMS. Network operators have between one and 50 such SMS centers, but there is only one platform type. Networks also tend to have other platforms such as value-added services platforms.

With MMS, on the other hand, there are several key platforms within the MMSE, including the MMS relay, the MMS message store, the MMS user database, and other platforms, including the existing platforms such as the SMSC, voice-mail platforms, and the like. There may be several of these, and they may be distributed as components in an open environment or integrated together in a single physical place.

The MMS infrastructure vendors have integrated the MMS server and relay into a single platform that they are calling the MMS center (MMSC).

Applications

In SMS, the vast majority of the total SMS traffic volume is accounted for by one application alone—simple person-to-person messaging in which people send messages such as "I'm bored" and "I'll be 5 mins late" from phone to phone. In the MMS world, however, Mobile Streams is predicting that initial applications will be application-to-phone, such as those sent from Internet sites and based on still images—for example, screen savers, mobile pictures, photos, postcards, and the like. In other words, the notion that MMS will be like SMS in that person-to-person messaging is the key application is naïve, because for the first couple of years before MMS reaches critical mass, most MMS messaging will involve a PC to either initiate or terminate the multimedia message.

User Behavior

One of the great advantages of SMS is that it is a very discreet communications medium that can be used "under the table" in meetings and at school, for example. MMS is almost the opposite of this in terms of user behavior insofar as people using a built-in or attached camera in many scenarios will draw attention to themselves. On the other hand, SMS–based enhanced messaging applications such as ringtones are also a means to draw attention to yourself, so perhaps the gulf in user behavior is not quite so great. This difference in user behavior is critical in educating users about MMS and is another factor that needs to be taken into account when it comes to MMS adoption rates.

Summary
MMS takes a lot of the winning features of SMS but improves and extends those capabilities with richer media. The MMS standards have been designed in a very elegant way to take the best of text and improve the rest. However, it is also very clear that SMS and MMS are very different in the way the services are set up, with MMS being far more complex, and also in the types of applications and devices used and the user behavior. Mobile Streams' is predicting that these differences will have an adverse effect on the rate of take up of MMS, as people will continue to use text in many communication scenarios.

3. Timescales

When a new service is introduced, there are a number of stages before it becomes established. Service developments for the MMS will include standardization, infrastructure development, network trials, contracts placed, network rollout, availability of terminals, application development, and so on. These stages for the MMS are shown in *Table 3*.

4. Milestones

Mobile Streams has developed a framework called "messaging milestones" to explain what steps are needed before multimedia messaging can become a success and how to maximize MMS volumes and revenues. The dates given in the parentheses are the approximate timelines for these milestones.

GPRS Networks Rolled Out (2001)
MMS is a 3GPP standard that can be run over GPRS or 3G network bearers. As such, this investment in underlying network technology is a critical one before MMS can be rolled out as a service using that bearer.

GPRS and MMS Billing Tariffs Sorted (2002)
It is important at this stage to start thinking about and investing in billing systems that can cope with MMS style services, including subscription services, service rather than volume-based billing, revenue-sharing arrangements with third parties, and the like, so that the new MMS–based services can be monetized.

First-Generation MMS Center Trials (2001, 2002)
Network operators purchase or trial the first-generation MMSC. As we have seen from the contracts awarded so far, there are a number of MMSC trials being implemented by network operators, mainly in continental Europe. Obviously, messaging infrastructure to manage MMS is crucial, since without it, MMS cannot be offered.

Business Partners Programme (2001)
It is important that network operators start at this stage to think about partners for MMS services. Consumer services will use the MMS synchronized multimedia intergrated language (SMIL) format, so a content creator with that capability would help. Corporate-oriented applications extending unified messaging or caller ID would constitute useful applications at this time too. A partner with a media bank for content and postcard and screen saver content could help here.

Initial MMS Terminals (2001, 2002)
New phones are needed before MMS can be used. As the initial MMS terminals become available, the first multime-

TABLE 3

Date	Milestone
2000/2001	Continuing GPRS, 3G, and MMS standardization with network architectures, terminal requirements, and detailed standards.
Early 2000	First MMS terminals were announced (Ericsson T68).
2000/2001	3G licenses for Phase 1 spectrum were awarded by governments around Europe and Asia.
2001/2002	GPRS networks are rolled out commercially around the world.
Mid-2001	Release 4 of 3GPP MMS specs frozen. WAP–MMS specs (based on pre-Rel-4 3GPP specs) frozen in WAP 2.0 suite. First MMS infrastructure contracts are awarded, and is it shipped for MMS trials and commercial services.
Late Q1 2002	First MMS terminal (Ericsson T68) is commercially available.
March 2002	Release 5 of 3GPP MMS specs frozen.
Mid-2002	M-Services Stage 2 devices, of which MMS is a compulsory part, start shipping.
2002	MMS infrastructure trials and contracts are placed in Europe, North America, Asia, etc.
Late 2002	Commercial volumes of MMS terminals begin shipping.
2003/2004	MMS will have reached critical mass in terms of the installed base of MMS capable terminals.

Source: Mobile Streams

dia messages are sent. The very first MMS terminals, technically speaking, are personal digital assistant (PDA)–type devices with GPRS equipment attached. The first MMS–capable devices, such as the Ericsson T68, will become available. At this stage, the network, software partners, and terminals are in place and services can start to be used, albeit on a small scale initially.

GPRS Roaming and MMSC Interworking for Inter-Network Message Termination (2002)

The addition of interworking between network operators that are competing in the same geographical market gives customers to all networks the opportunity to use MMS in the same way as they do voice. Just as they can make a voice call to each other's phones, so too can they send multimedia messages to each other. Enabling this capability can rapidly increase the number of available messaging destinations, thereby increasing the value and use of MMS.

MMS Terminals and Service for Prepay

The next increase in MMS traffic volumes will be caused by the introduction of MMS terminals on prepay packages. Network operators face an issue in that they want to subsidize MMS terminals to attract usage but do not know what return on that subsidy they will get on a prepay tariff with no fixed contractual length. The youth market was the early adopter of SMS and is likely to be interested in MMS, yet the terminals are likely to be relatively sophisticated, and therefore expensive.

Range of Interoperable Terminals (2002, 2003)

Obviously, it is important that a choice of different terminals is available to potential MMS users from a range of terminal manufacturers. All of these terminals must closely conform to the 3GPP and WAP–MMS specifications and interwork with handsets and MMS infrastructure from a range of different vendors.

End-User Awareness and Education

Obviously, user awareness and education are important, although MMS is expected to grow very quickly in this regard, since it is instantly clear from using an MMS–capable device that the messaging richness is dramatically improved compared to SMS and text messaging.

As such, there are various steps that mobile network operators can and should take to spur the initial development of MMS usage. Each of these steps is complementary and useful in making MMS a success. All of the stages are crucial to the overall success of MMS. It is not until these steps are in place that MMS can start to really succeed. These are the initial factors—networks, services, terminals, billing—before the MMS traffic can really start to take off.

5. Technical Features

The enablers for multimedia messaging have their roots in the changes that are taking place at all levels within the mobile Internet. Enabling bearers such as GPRS, EDGE, and 3G are becoming available. Enabling technologies such as Bluetooth, WAP, MexE, and SyncML are all initiatives that support this new direction toward the mobile Internet. We are also seeing new categories of multifunctional devices, such as MP3 phones, being implemented.

The 3GPP has been studying the limitations of SMS, and the MMS is the answer to the new messaging market requirements. Infrastructure vendors see multimedia messaging as an evolution in mobile messaging as the user moves from SMS to EMS to MMS. The MMS is, according to the 3GPP standards, "a new service which has no direct equivalent in the previous ETSI/GSM world or in the fixed network world." Let's introduce the features of this innovative new service:

1. MMS is a presentation layer for e-mail such that it uses underlying standard e-mail protocols but presents that e-mail in a compelling design.
2. MMS will enable messages to be sent and received using lots of different media formats and types, including text, images, audio, and video.
3. As new more advanced media become available, more content-rich applications and services can be offered using the MMS service environment.
4. The MMS introduces new messaging platforms to mobile networks to enable MMS. These new platforms have been designed to interact with legacy mobile platforms such as SMSCs. The new platforms include MMS relay(s) and MMS user databases.
5. MMS is like SMS—a non–real-time service. A relay platform routes multimedia messages to MMS servers.
6. The MMS is designed to be future-proof. As mobile networks evolve and new media become available, the aim is to make the standards as backward- and forward-compatible as possible.
7. Access to MMS services should be independent of access point—multimedia messages should be accessible through 3G and 2G mobile networks, fixed networks, the Internet, etc. This is where common message stores will be an important enabling technology. To facilitate interoperability and universal messaging access, MMS will comply with the virtual home environment (VHE). VHE is a 3G concept that simply lets customers have seamless access with a common look and feel to their services from home, office, or on the move and in any city as if they were at home. The VHE permits the user to manage his or her services (including non–real-time multimedia messaging handling) via a generic user profile, permitting, for example, all different types of messaging to be presented to the user in a unified and consistent manner. (See www.mobileVHE.com from Mobile Streams for more information.)
8. The concept of a user profile has been included in the MMS standards. This user profile is stored in the mobile network and is user defined and managed via the Internet, and determines which multimedia messages are downloaded immediately to the user and which are left on the server for later collection. The user may also choose to receive notifications of certain multimedia message types.

   User profiles are important for several reasons, including the fact that different MMS devices will support different levels of capability. The user profile will also allow the user to control access and accessibility for multimedia messages that will be mobile originated and mobile terminated. The user profile will include filtering rules and routing tables for unsolicited or undesirable messages.

9. Although MMS is being standardized by the 3GPP, in fact MMS services can be offered on any IP–based air interface, such as GPRS (so-called 2.5G) or EDGE networks. The MMS is a service environment that interworks with and allows multimedia messages to be transmitted across IP–based mobile networks. (See www.mobileGPRS.com, www.mobilEnhanced Data.com, and www.mobile3G.com and the "Success 4 GPRS" and "Yes 2 3G" reports for more information on these bearer networks.)
10. Content scaling (e.g., downscaling of images) and content transformation (e.g., conversion of one audio format into another) have been considered in the MMS standardization process.

In summary, the MMS provides an intelligent environment for multimedia mobile messaging.

6. Architecture Elements

This next section gives an overview into the technical implementation of MMS. It quotes heavily from the 3GPP specifications, in particular 3G TS 23.140, a functional description of MMS.

The multimedia messaging architecture has a number of key elements that have been defined and incorporated into a MMSE. The concept of the "service environment" has been used in the past, for example in intelligent network (IN) solutions with the CAMEL service environment (CSE).

The MMS architecture contains several key platforms that interwork with each other to provide the MMS service. The key elements defined by the 3GPP are as follows:

- MMS relay
- MMS server (or servers)
- MMS message store (or stores)
- MMS user agent
- MMS user databases

It is important to note that MMS infrastructure vendors are integrating MMS platforms, such as an MMS relay and MMS server, and labelling them as an MMSC.

To understand these platforms, think about the internal organs of human beings. The MMS relay is like the heart of the service, which pumps content around the body. The MMS server(s) is like the other organs that perform other related functions. The MMS message store(s) store related fluids that get passed around. The MMS user databases contain the profiles that decide how the multimedia messages are transmitted and displayed. The MMS user agent is a like the eyes—it is the external place through which the processed information is displayed and viewed.

This architecture allows multimedia access to all types of different information with a range of servers providing access to new and legacy services. This allows operators to consolidate access to multiple applications from a single architecture.

Figure 1 shows that multimedia messaging (MMS) may encompass many different network types that can be connected by standard IP messaging formats such as SMTP, MIME, etc. This approach enables messaging in 2G and 3G mobile networks to be compatible with Internet messaging systems.

Having discussed the criteria for selecting an MMS vendor, the MMS infrastructure vendors will now be reviewed. The profiled companies are ADC Enhanced Services, CMG, Comverse, Ericsson, Logica, Materna, Motorola, Nokia, Openwave Systems, SchlumbergerSema, Siemens, Tecnomen, and Unisys (see *Table 4*).

FIGURE 1

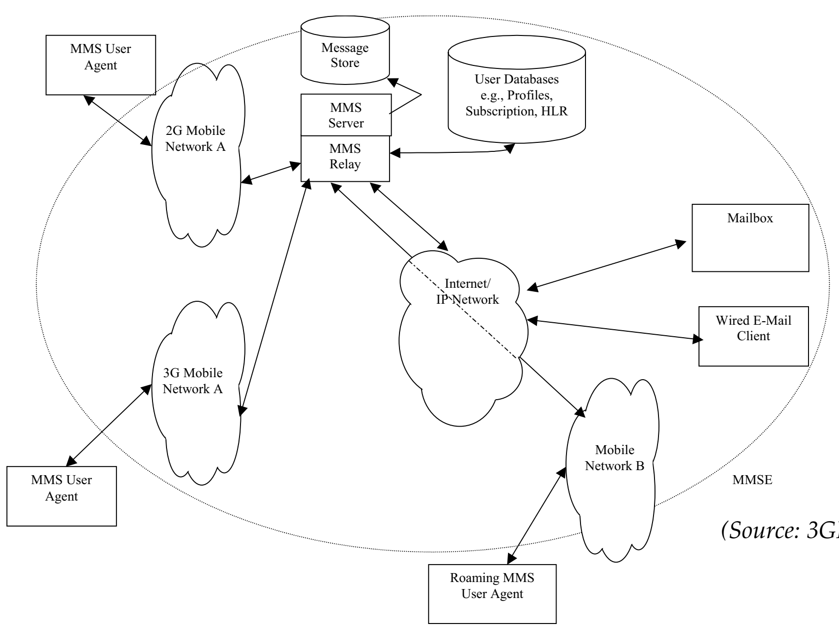

(Source: 3GPP)

7. MMS Applications

Mobile Streams believes that the key services and applications based on MMS will be related to mobile entertainment and still images in a person-to-person environment.

MMS Success = Entertainment + Still Images + Person to Person

Let's look at this in more detail:

Mobile Entertainment

Mobile entertainment services are key for MMS. They account for the majority of i-Mode traffic, for example, on the NTT DoCoMo network and are a more popular application than mobile messaging as a category of non-voice services. There are also likely to be substantial revenue premiums attached to mobile entertainment services—it is possible to charge a substantial premium for ring tones and picture messages in the EMS world, and this trend should be and is likely to be enhanced and maintained in the MMS world.

Person to Person

About 90 percent of SMS traffic volumes are still accounted for by simple person-to-person messaging—people picking up their mobile phones and sending each other messages such as "Hello how are you?" and "I'll be 5 mins late." In such an environment, SMS is a medium and the end users self-create most of the content themselves.

In MMS, once an installed base of MMS users have been built up, person-to-person communications will also be the majority of the traffic volumes, although mobile phone users are likely to depict their feelings in visual rather than textual format. With MMS, they may compose a picture that they created themselves on the MMS terminal or a photograph that they took of themselves or something else or one they found on a Internet site and sent to their friend after adding text to the picture. One to one communications from one person to another are likely to still be the overwhelming application usage in the MMS world, with the added visual element.

While simple person-to-person messaging will start off slowly before MMS is widely available in the mass market, we still believe that sharing the moment will be a popular service for MMS.

Still Images

Still images such as screensavers, photographs, pictures, letters, postcards, greeting cards, presentations, and static Internet pages can be sent and received over mobile networks just as they are across fixed telephone networks.

Two variables affect the usability of such applications—bandwidth and time—and they are inversely related. The faster the bandwidth, the less time is needed to transmit images, and vice versa. This is the reason why the transmission of image-based rather than textual information has not been a popular non-voice mobile application until now—it takes too long given the slow data transmission speeds that were available prior to the introduction of GPRS mobile packet data.

Once captured, images can then be sent directly to Internet sites, allowing near–real-time desktop publishing. The size

TABLE 4

MMS Vendor	End to End	Vision/ Strategy	Key Platform	SMS	WAP	Voice	UM	IP	Contract	Partners
ADC	Low	Med.	UM	Med.	Low	Med.	Med.	Low	Low	N/A
CMG		Med.	SMS	High	High	Low	High	Med.	High	Med.
Comverse	Med.	Med.	UM	High	Med.	High	High	Low	Med.	Med.
Ericsson	High	Med.	SMS	Med.	High	Med.	Med.	Med.	High	Med.
Logica	Low	Med.	SMS	High	Med.	Low	Low	Low	Low	Med.
Materna	Med.	Med.	New Platform	Low	High	Low	Med.	Low	Low	High
Motorola	High	High	SMS	Med.	Med.	Low	Low	Med.	Low	Med.
Nokia	High	Med.	New platform	Med.	High	Low	Low	Med.	Med.	Low
Openwave	Med.	Med.	E-Mail	Low	High	Med.	Med.	High	Low	Med.
Sema	Low	Med.	UM/SMS	High	Low	Med.	Med.	Low	Low	Med.
Tecnomen	Low	Med.	UM	Low	Low	High	High	Low	Low	Med.
Unisys	Low	Low	UM	Low	Low	Low	High	Low	Low	Low

Source: Mobile Streams

Note: A rating of "High" denotes a better ranking than "Med." and so on.

of the file for a picture depends on the resolution and type of compression. Typically, each picture is between 50K and 100K in the Joint Photographic Experts Group (JPEG) format. This can be transmitted quickly using mobile packet data networks, especially 3G networks.

Still image transmission is a much-touted application for lower packet data services such as GPRS and beyond. Many people see still images as a killer application for GPRS and 3G.

While moving images, video, and the like represent a service that is compelling for sports replays, news clips, and the like, it is thought that it will be a relatively small market size for MMS compared to the potential for transmitting still images.

8. MMS Phones

Ericsson T68

The Ericsson T68 will be the industry's first phone to support multimedia messaging (MMS) standards. The Ericsson T68 will also be the first mobile phone from Ericsson with MMS, a color display, digital imaging, and audio capabilities. It is a radical phone in other ways too, as this review will show. The T68 has been developed for GSM 900/1800/1900 and GPRS networks and will be shipping in volume in late Q1 2002.

The first impression of the T68 is that it is very small—fitting easily into the palm of your hand. It is immediately noticeable too that the T68 is also very light. The T68 measures 101x48x19 mm and weighs 85 grams. The T68 has a large eight-row, 256-color display that is very clear and easy to read. The backlit screen is very bright when the phone is in use, but to conserve battery life, the standby mode turns the phone into an almost monochrome status.

The next thing—after the size, weight, and screen—that the user notices is the joystick navigation. A small joystick is centered below the phone screen and is pressed to issue a command or moved up and down to scroll between menu options. There is also a scroll button like a "jog bar" on the side of the phone that can be used, for example, to find out the phone status (battery life and the like). The advanced state of the physical design of the T68 continues to impress with an integrated antenna and even the back of the phone is curvaceous and well thought out without being overdesigned like some of Nokia's newest contrivances.

When you press the joystick, you get a visual menu with icons for different things, such as messages appearing on the screen. The software man/machine interface (MMI) is as surprisingly intuitive as the hardware design. Ericsson has finally and belatedly realized (along with the likes of Motorola and Siemens) that intuitive menus are essential for the usability of the device.

For example, the Messages menu lists Voice-Mail, SMS, EMS, MMS, E-Mail, Chat, Area Information, and Options, and the Fun & Games menu features Themes, My Pictures, Draw Pictures, My Sounds, Composer, Sound Recorder, Games, and Plug-In Camera. The phone supports all kinds of messaging (SMS, EMS, MMS) and even supports (unofficially) picture messages based on the Nokia Smart Messaging format! These cool features mean that you can personalize the phone with an entire theme (such as Microsoft Plus! on PCs) with wallpapers, sounds, ring tones, and the like. I alternated between the psychedelic and penguin modes depending on the time of the day! The sound-recorder feature works as its name suggests—for example, I recorded the sound our credit-card machine makes!

The T68 does not have an integrated camera but can support an add-on; however, this is not an inhibitor, because the phone can use many services that do not rely on taking pictures with a camera. For example, the T68 supports downloading 60x80 color images (MMS pictures). You can also "Add Page" for sequences of multimedia pictures and control page timing—the rate at which the pictures refresh as a multimedia presentation is played. These "screen shows" and wallpaper backgrounds give more personalization options for the phone user.

Within the MMS sub-menu, you can get a feel for the similarities between SMS and MMS. The options are Validity Period, Read Reply, Delivery Report, Auto Delete, Auto Download, Service Center, and WAP Profile. Many of these features, such as delivery confirmation, will be familiar to SMS users. When you write a new MMS message, you select a picture from the on-board pictures (which can be topped up over the air), and then add text to the picture and send it to a mobile number or e-mail address.

The T68 is a close to perfect execution of how an MMS phone should be. It is incredibly difficult to fault any aspect of the phone. The only thing you can really say is that the price of the phone is not yet clear, although it is expected to be priced aggressively. The only other thing I could come up with as a criticism was that the metallic gold finish on the device appeared very flimsy and a little tacky up close, although no-one can mistake this phone for anything other than a very, very serious piece of cool kit. The MMS device supports a picture plus text, and not animated messages or other MMS features, although with slide shows and those features, most users will not know or care about this.

Ericsson is working to ensure that the T68 is interoperable with various MMSCs from Nokia, Comverse, and CMG, as well as Ericsson, of course. Ericsson is not working with Logica for terminal interoperability at this stage, which is a surprise. Ericsson has shipped the T68 for trials with Europolitan-Vodafone in Sweden, for example.

Ericsson has set the bar very, very high with this mature implementation of the first MMS phone. Picking up the phone is really one of those pivotal moments in your life, such as when you browsed the Internet for the first time, when you realize just how great the difference between messaging now and tomorrow will soon be. Because this revelation comes as a gadget in a phone package, it has an even stronger impact. The T68 really brings MMS to life.

Ericsson has found what seems to be the perfect form factor for MMS—a large enough screen size, but not too large, with great hardware and software navigation. Because Nokia and now Ericsson have developed intuitive user interfaces for MMS, this is not likely to be as big an issue as it was in the SMS world before predictive text input was introduced. It is perhaps no surprise that "T68" rhymes with "great," because the Ericsson T68 multimedia phone really is very, very great!

The product I tested is a software upgrade version, which Ericsson will introduce during Q1 of 2002. The launch date was delayed because, according to Ericsson, "MMS development has not been as rapid as we expected which means we haven't been able to test the product as thoroughly against the various server vendors. Therefore, we cannot introduce MMS just yet." Ericsson will ship T68 without MMS during 2001 and introduce MMS in a later release toward the end of Q1. It will be possible to load MMS the first version of T68.

9. MMS Statistics

Global MMS Market Forecasts 2002–2008

Mobile Streams is predicting that MMS will start to grow in 2002 but will not really take off until 2004. We expect MMS to continue growing in volume for many years to come beyond 2003, until at least 2009 (see *Table 5*).

TABLE 5

End of Year (December Figures)	Monthly Total Global MMS Market Size	Monthly Total Global SMS Market Size
2002	$25 Million	$62.4 Billion
2003	$350 Million	$80 Billion
2004	$2.5 Billion	$82.1 Billion
2005	$10 Billion	$79 Billion
2006	$20 Billion	$59 Billion
2007	$50 Billion	$40 Billion
2008	$65 Billion	$26 Billion

Source: Mobile Streams

These figures denote the monthly figure for December of that year, *exclude* all other message types (such as SMS and EMS) and show only MMS, are approximate, and exclude Japan.

The Evolution of Softswitch Architecture

Walter Buga

Co-Founder, Chief Executive Officer, and Chief Technology Officer

IP APPS

Director, Technical Advisory Committee

International Softswitch Consortium

Introduction

The key issues discussed in this paper are softswitch architecture evolution, visions and realities, and challenges and benefits. The paper considers these challenges from a greenfield service-provider perspective.

Softswitch Architecture Evolution

The traditional public switched telephone network (PSTN) is based upon circuit-switching technology, while the new softswitch-based networks will leverage packet-switching technologies. This new architecture will simplify service creation and will integrate today's voice and data networks into one converged one that can provide both services. However, this change will not occur overnight, as there is more than $350 billion of legacy equipment installed in the PSTN today. Thus, this transformation will be evolutionary rather than revolutionary. In the section below, the stages of such an evolution are described.

PSTN Model

In the most basic and simplified view, the PSTN is composed of three layers:

- **Switching fabric:** connected to the transport layer via lines to end users and trunks to networks
- **Signaling and control:** provides call-control functionalities, including call establishment and teardown
- **Applications and services:** available to the end user, such as Class-5 features (call waiting, call forwarding, etc.) and advanced services (toll free, calling cards)

Typically, these three functionalities were implemented in "one box" as a Class-5 or Class-4 switch (see *Figure 1*).

Early Softswitches

The trend toward more open and programmable switches existed for a number of years, but their market penetration was very limited. Early adopters attempted to separate switching fabric from other functionalities via an open (but proprietary) interface/application programming interface (API) (see *Figure 2*). These so-called programmable switches were used for niche applications, such as callbacks, calling cards, etc.

Distributed Switching Fabric

Using packet-based networks—Internet protocol (IP) and asynchronous transfer mode (ATM)—to transport voice traffic, instead of the traditional time division multiplex (TDM) technologies used in the PSTN, allowed for switching fabric distribution. In these new networks, packetization of voice starts at a user device (IP telephony [IPTel], integrated access device [IAD], etc.), and then packetized voice is transported/switched to the "end" of the packet network. Finally, it is converted back to the TDM format before exiting to the PSTN. *Figure 3* depicts this architecture.

Open Architecture

The trend to break the old paradigm, in which a single vendor supplied everything—software, hardware, and applications—in one proprietary box, was gaining more traction because customers did not want to be locked in to their vendor with no room for innovation.

In mid-1999 the International Softswitch Consortium (ISC) (a nonprofit organization) was formed to promote an open architecture and multivendor interoperability for next-generation voice/video/data solutions. In less than two years, its membership grew to more than 180 system suppliers and service providers. This reflects the industry momentum toward this new architecture.

Open standards enable innovation and reduce costs, and customers are free to choose best-in-class products to build their network. Much work needs to be done before that vision can be realized, as many interfaces still need to defined and standardized.

Figure 4 depicts the target open architecture to be realized. This is only a functional view, and some vendors may elect to include various elements of this architecture in one box,

FIGURE 1

PSTN Model Switch

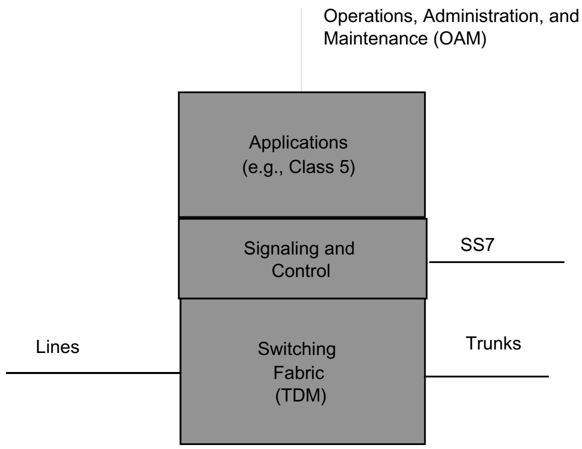

FIGURE 2

Early "Softswitch" Architecture

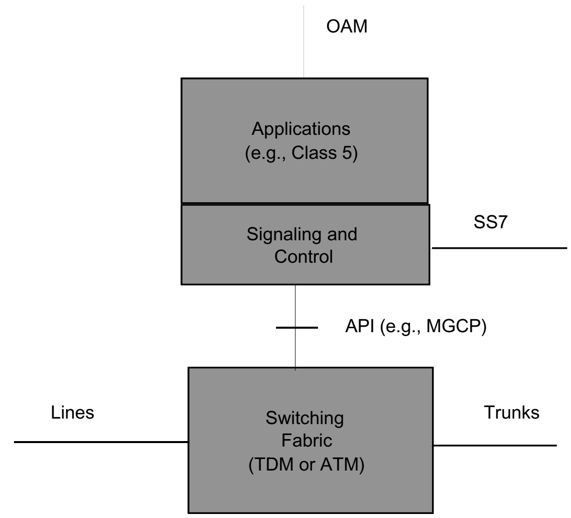

while others may wish to focus on a very specific element(s) of this architecture.

The Next-Generation Converged Network

Some next-generation service providers want to build voice services on a softswitch platform that can be deployed on a fully converged broadband and IP–centric network with quality of service (QoS)–managed virtual pipes using multiprotocol label switching (MPLS). This arrangement allows for an easy deployment/add-on of new advanced services and integration with the advanced service-creation platform. In this arrangement, voice applications can be considered one of many application service provider (ASP)–type offerings.

The architecture for such a network is a packetized network built upon a switching fabric that could be either IP or ATM. Everything else in the network is built on top of this switching fabric and integrated through open/standard APIs.

Thus, the greenfield service provider could start with a basic packetized network—no legacy elements are included if they can be avoided. To add a new customer, an IP–based phone or IAD would need to be added on the front-edge customer-premises equipment (CPE) of the network to convert voice to packetized form (including compression). This packetized voice traffic could be terminated "on-net" on another IP phone or IAD. However, for the "off-net" calls (the ones that need to exit to the PSTN), a media gateway (MG) would need to be deployed to convert this packetized traffic back to TDM format. Providers that do not have signaling system 7 (SS7) networks of their own would to be able to connect to providers of SS7 networks to provide call-control and signaling functions for voice traffic to/from the PSTN.

FIGURE 3

Packet-Based Network Architecture

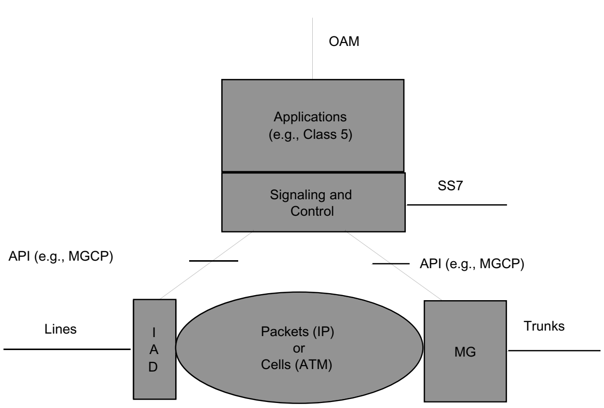

FIGURE 4

Open Architecture

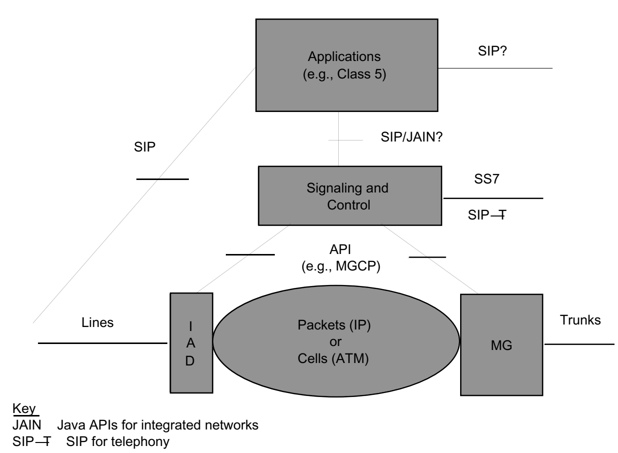

An MG control protocol (MGCP), or MG control (MEGACO) (when available) could be used to control MGs (including IADs). Access to applications could be provided via a session initiation protocol (SIP). Once this IP network is in place, other applications and content, such as video and audio, can be easily added using a standard IP suite or other standard interfaces.

Figure 5 depicts an example of such a converged network.

The Vision

Unmanaged and Free to All?

It can be assumed that, in the future, all devices and applications will be IP enabled and connected to the ubiquitous IP network. Users and/or applications will communicate/interact via virtual, secure, and QoS–controlled/managed IP channels. Broadband access to the core IP network will be provided via a variety of economically justifiable physical media, including, wire, wireless, cable, optical, etc.

The IP network is location agnostic, as it is concerned with only the IP address, the bandwidth, and the QoS available for a specific connection(s). Assume that a telco has a data center, applications, and content connected to this IP network, while a residential customer has a phone and a browser connected to the same IP network. The difference between the telco and the residential customer is that the telco has a bigger pipe. In both cases, all the applications reside on the edges of the network, including users, telcos, ASPs, access devices, and applications (see *Figure 6*).

The number of users, devices, and applications connected to the Internet is constantly growing, enabling more and more choices. A user must have access to a small set of applications available on the Net. In the early years of the Internet, search engines helped to navigate the Net. Later, vertical portals, such as Pet.com, and horizontal portals, such as Yahoo (evolved from the most popular search engines), helped to navigate the Net. Such issues as QoS and network security were ignored or left to the end user to deal with.

Service-Provider Vision

What must the service provider do to implement the vision of a manageable and secure Internet?

The first necessity is content and applications, which someone must produce. Another requirement is service creation. The provider should be able to customize and provide an application, or a set of applications, for its customers. The provider may wish to create its own applications or to have access to others that are creating applications for the Internet. Then the provider must deliver and manage the service by provisioning, configuring, maintaining, and charging for it. Because voice applications are part of the service portfolio, control and management of QoS is a must.

The service-provider vision must go beyond the typical Internet, where every user can access everything out there. To realize that vision, service providers must implement middleware functions—such as security, directories, policies, etc.—and back-office systems to provide the necessary support.

The Challenge

Realizing this vision requires surmounting some obstacles.

One obstacle is product maturity, particularly in a voice/softswitch space. Until recently, a provider could choose

FIGURE 5

Converged Architecture

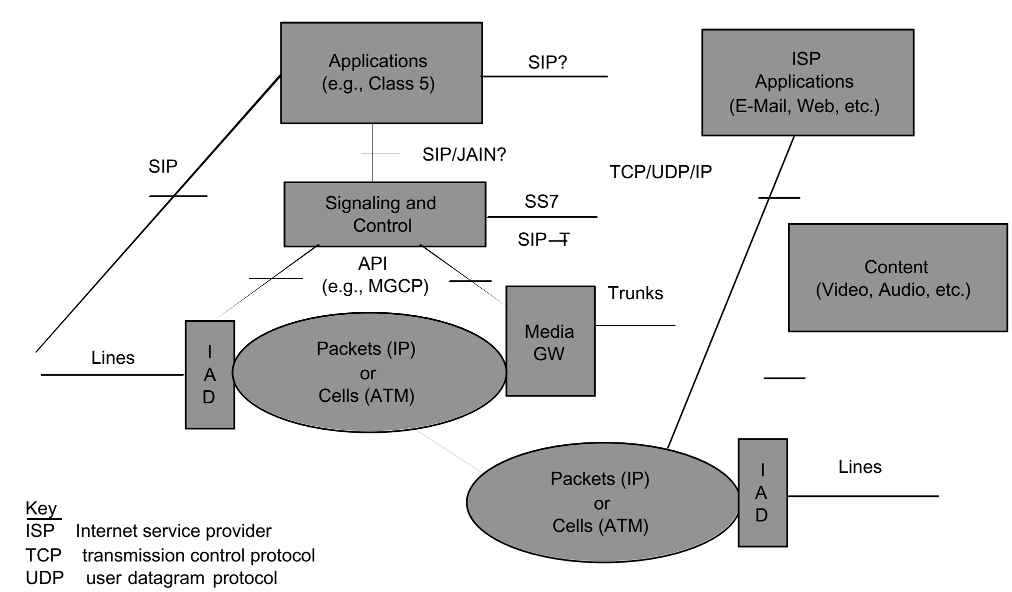

FIGURE 6

Vision

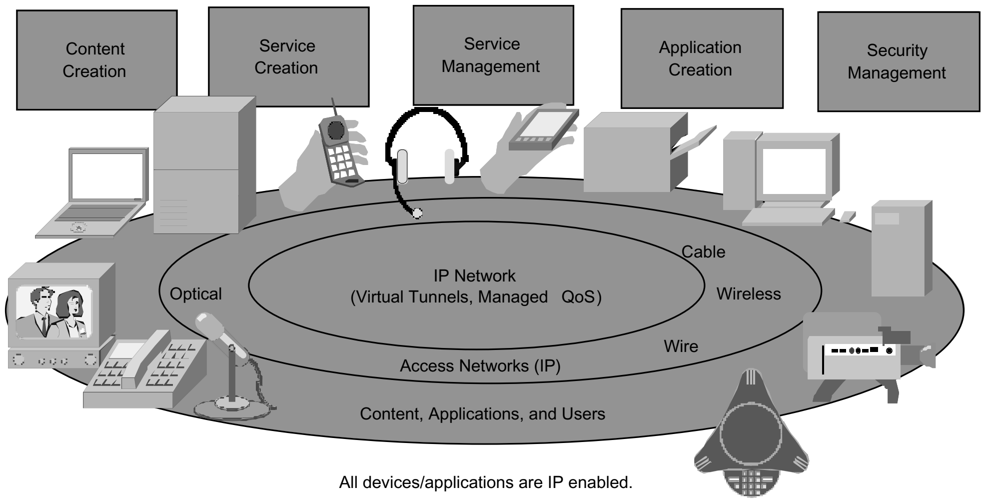

from only a few products. This picture is improving, however, and new vendors are bringing appropriate products to market. Product-release dates and schedule delays are other concerns. They create numerous problems as providers try to design a network that will implement the vision and meet target schedules, all while dealing with multiple vendors. Price is a key obstacle. Even if the provider offers SIP, customers will not accept the offering unless the price is right. But the provider must generate revenue in accordance with the business plan to survive in today's market.

Deployment

What is the solution? Deploying the best available solution and evolving it as soon as possible is the only way to overcome these obstacles. This approach will accomplish the vision.

Deployment of new technologies to existing networks carries some risks, but deployment of cutting-edge technologies and services in greenfield networks involves a very high risk—from scheduling problems and interoperability to the actual operations.

As an example: The greenfield voice service provider must be concerned with at least the following:

- Getting all common language location identifier (CLLI) codes in place
- Getting numbering plan area (NPA)–NXX numbers (pooling, lotteries, time limits, cumbersome processes, etc.)
- Getting SS7 connectivity
- Getting 411 services
- Obtaining Communications Assistance for Law Enforcement Act (CALEA) compliance or waiver
- Bridging a "culture" gap between the PSTN and the Internet

It is difficult to address the above because the process is cumbersome, random (e.g., lotteries), and not well defined, and the solution must then be implemented within a specific time frame or the process will need to be restarted. And perhaps, overcoming the cultural gap between the PSTN and the Internet is the biggest challenge.

However, the advantages are clear:

- Capital costs of a softswitch are dramatically lower than with a traditional circuit switch.
- Softswitch technology scales better, making it possible to employ a "pay as you grow" approach.
- IP switching also enables extremely efficient use of bandwidth
- The softswitch fits well in an IP network and enables an easy transition to IP services.

Conclusions

This discussion has presented a vision of an IP converged network that can offer a portfolio of new IP–based services and applications, including softswitch-based voice.

Obstacles to implementing the vision include product maturity, scheduling problems, operational challenges, price, and interoperability. Deploying the best available solution and evolving the network as new possibilities emerge can overcome those obstacles and lead the provider toward the goal: the fully realized vision of a fully converged and manageable network and services.

The Past and Future of Wireless

Martin Cooper
Chairman and Chief Executive Officer
ArrayComm

Business people in the United States are transfixed on the war on terrorism, the recession we've just entered, and the inevitable but elusive recovery. But the telecom industry has been in recession for some time. In fact, the telecom industry is in meltdown. Just look at the continuing layoffs, shrinking valuations, once respectable companies in a cash crunch, and some companies in deep trouble, and the only discussion, once again, is not "will there be a recovery this year," but rather, "will there be a recovery, at all, in 2003."

I respectfully suggest that there will not be recovery for most of the industry. Rather, in the long term, we will see a rebirth—a new telecom industry rising above the residue of the meltdown. Our industry has finally, after 120 years (yes, that's how long ago Alexander Graham Bell invented the telephone), started to grow up—started to escape from the burden of its monopoly heritage.

Things *will* get better, of course, but, in the short term, we will only encounter an illusion of recovery. The real recovery, the rebirth, will involve fundamental changes in our industry. Our industry will change because it has to change

The wireless segment is in just as much trouble as rest of the industry, but that trouble is largely masked by the continuing growth of subscriber usage. Wireless—cellular—is, today, an unfulfilled vision. Performance is marginal-to-poor everywhere. Infrastructure deployment is not keeping up with traffic growth, but it's difficult to find the capital to deploy more infrastructure when revenues grow only upon the addition of *new* subscribers. We are in recession, and there are just not a lot of new subscribers.

How will carriers get revenues going up again? The carriers have figured out that voice is just one data application and that the solution to the growth need is data applications. Data is now the Holy Grail—the brass ring.

But there's a serious obstacle. Our industry, all telecom but particularly wireless, has inherited some bad habits out of our monopoly legacy. Let me use the creation of cellular as an example. Let's start with some history:

The essence of our vision of cellular, when we developed it in the late 1960s and 1970s, was the unleashing of the consumer from the wires that chained them to their homes, offices, and desks. We were going to replace the wired telephone as the communications tool of an increasingly mobile population. We made good progress from the birth of cellular in 1983, but then "digital" wireless technology came on the scene and somehow there was a change in attitude. We lost our vision. The dream was replaced by hype and exaggeration totally unrelated to consumer benefit.

Digital technology did offer advantages to wireless operators. There was an increase in capacity—more subscribers could share the increasingly costly radio channels—but the only real advantage to the consumer was a valuable increase in battery life. There was a trade-off—voice fidelity wasn't quite as good as in well-designed analog systems, and phones cost more. But there was a net gain in subscriber value.

Then the "air-interface war" started. It became a religious war, unhindered by facts, with billions of dollars spent to change things—but often for changes based upon hype and exaggeration, on meaningless acronyms and technology-for-the-sake-of-technology with no real contribution to the original vision. We got into the habit of talking about technology to the exclusion of what should have been our topic—the consumers and their needs.

Third-generation cellular, 3G, is the culmination of that era of hype. For a while, we actually believed the 3G hype, and some still do. But the engineering and economic facts have been evident for years.

The Hype: A huge increase in subscriber capacity and delivery of data at speeds of 2 megabits per seconds (Mbps).
The Reality: A slight increase in capacity (we've already extracted most of the potential improvement in capacity from air-interface technology), and data speeds of just 64 to 100 kilobits per second (kbps).

The Hype: Delivery of multimedia.
The Reality: We're still trying to figure out what applications are practica,l but *multimedia*, in the traditional sense, is just not economically or technically possible.

The Hype: Internet on the move for *everybody.*
The Reality: The Internet has the potential to make practical a host of new applications that improve productivity, convenience, and safety as well as entertain and educate—but do you really believe that this can happen through a 4 centimeter window at 64 kbps on a device designed for voice over a system that was also designed for voice? NOT A CHANCE!

3G will happen! But it will serve specific markets and applications. Voice will continue to be major factor in cellular systems, but there will be a new class of systems that are optimized for Internet delivery and are unburdened by the need for "universality" that adds so much complexity and cost to 3G. 3G will happen when the focus returns to the customer—where it is in every other industry, not burdened by the monopoly heritage of the telecommunications industry that requires the carrier to "own the customer." 3G will happen when new technology is embraced that drives capacity and performance to much higher levels than 2G or 2.5G—technology that exist today but hasn't yet become a part of the 3G story. For example, there are smart antennas deployed today in more than 90,000 cellular base stations with proven capacity increases of 10 times or more, all at *lower* cost—but the traditional cellular industry has resisted this crucial innovation. The industry *must*, and will, open its mind to such new technology, and those companies that do will be the new leaders in the reborn telecom world.

i-Mode gives us an interesting peek into the future. Applications on i-Mode are not very profound; some are even trivial. The data rate is only 10 kbps. But i-Mode represents the first crack in the "walled garden;" the i-BURSTMode carrier, NTT DOCOMO, invites lots of creative entrepreneurs to find and exploit market niches. i-Mode is just a crack in the walled garden because NTT DOCOM still bills the customer (remember the legacy of "owning the customer"?) But just imagine what these entrepreneurs could do with, say, a megabit per second that is delivered at *really* low cost.

No one system will fulfill all of our wireless future. When the telecommunications industry grows up—and that process is happening now—we'll start with a service or an application, or a device that brings value to our customer, and then we'll invent the systems that optimally deliver those services to the customer. 3G will be one of those systems, but look for others—they will be there.

The wireless industry is in its infancy, but it's about to blossom out. Engines of that blossoming will be a new business attitude that focuses on the customer—the open platform—and new technology, such as adaptive smart antennas, that multiplies the use of the spectrum and new systems, such as ArrayComm's i-BURST, that are optimized to serve specific, previously unserved markets. The blossoming won't happen quickly—disruptive changes seldom do—but we'll see major changes during the next 10 years.

And what companies will survive? The ones that break out of monopoly thinking, that start with the customer and embrace, that create technologies that serve the customer best, that carve out market niches (perhaps huge ones), that serve classes of customers better than others do, in short, those companies that adopt strategies that release them from the "copper cage."

Notwithstanding the long-term nature of my predictions, there will be short-term opportunities in the companies that adopt new competitive approaches that serve lots of markets and that offer content that is meaningful—opportunities in technologies that really perform and in services that make people's lives better. You'll have to figure out who the winners and losers are. But things will change—don't say I didn't warn you.

This paper was presented at the Economist Conferences Seventh Roundtable on Asia-Pacific Telecommunications in Hong Kong on November 26, 2001.

Then and Now: A Perspective on Recent Changes in the Telecommunications Industry

Edward Cornet
Managing Partner
Booz Allen & Hamilton

The telecommunications and Internet industry has changed dramatically over the past decade, but never have changes been as abrupt and unsettling as those in the past two years. In this paper, I'd like to briefly review where the industry was then and where it is now, discuss some of the root causes of recent changes, and explore what the survivors should be doing going forward.

Where Was the Industry Two Years Ago?

Stock prices were high and revenues were growing. Investments were being made in new competitors. Infrastructure was being built at an astounding rate, and established companies were bold with new acquisitions to expand their footprints. The industry was enjoying a time of "easy money" for just about any good idea. Internal projects for new product development and market research were easily funded.

During the previous four years, a thousand new companies had started up to take advantage of the Internet growth spurt and the less regulated telecommunications market in the United States. Competitive and incumbent local-exchange carriers (CLECs and ILECs) as well as long-haul carriers were investing in infrastructure to handle the projected explosive growth in communications and Internet demand. Service providers were multiplying with special emphasis on non-facilities–based services and application services.

Public companies had "alliance capital" created by their strong stock price growth driven by their promises of future revenues. The inflated stock prices gave companies increased equity to expand their businesses. Banks were willing to increase debt ceilings as stock prices rose, even though revenue growth was speculative. As new infrastructure was built, deeper depreciation wells were dug in the balance sheet.

Some of the alliance capital was used for acquisitions to expand market footprint or add broader value-chain service capabilities. Some alliances had no apparent strategic objective, but rather seemed to be more of an exercise of management power. Even alliances with strong strategic rationale suffered from little or no serious post merger integration activity to streamline costs or capture value. Management was allowed to get sloppy as enthusiasm for future growth became widespread.

Businesses capitalized on some exciting technology trends. Voice/data convergence, voice over Internet protocol (VoIP), seamless open networks, and video streaming promised to unlock hidden demand. Dense wavelength division multiplexing (DWDM) and Internet protocol (IP) technologies would drive the cost per byte to zero. Advanced wireline technologies, multiprotocol label switching (MPLS), asynchronous transfer mode (ATM), IPv6 protocols over cable, fiber, or digital subscriber line (DSL) would enable broadband everywhere. Long-distance voice and international Internet access would become affordable for all. Third-generation (3G) wireless networks, code division multiple access (CDMA)2000, and wideband CDMA (WCDMA) were on the horizon. Users would be on line everywhere, all the time. Fixed/wireless convergence would create a total access world. Investment bankers and others were willing to fund just about any good idea tied to an advanced technology.

The concept of access to anything, anywhere, anytime seemed possible—and right around the corner. The big questions were: how quickly were we going to get there? And if you hadn't started already, how could you catch up?

What Happened?

The number of competitors grew rapidly, diluting the existing demand. Competitive pressures drove down prices. Thus revenues failed to materialize as promised. Meanwhile costs continued to grow. Positive cash flow was delayed. Additional investment was required, but investor confidence was lost. Start-ups were unable to secure additional capital to fund operations, and bankruptcy followed. Many non-infrastructure–based companies folded entirely. New infrastructure companies were restructured with their assets valued at cents on the dollar. Even the larger, established companies with positive cash flow were affected—forced to stop further investment and cut operating costs to meet shareholder demands. The bubble had burst.

What Were the Causal Factors behind the Collapse?

1. *The most common problems were incomplete business models.* In the dot-com world, the focus was the concentration of "eyeballs" on the Web site or page hits that generated the hype versus real customers or revenue streams. The survivors had models based on fundamentals. For example, Amazon.com's foundation as a bookseller is based on a very sound financial principal. It was competing against local bookstores. The Amazon.com principle was to significantly reduce the costs to the reader. They recognized that about 35 percent of the cost of the retail price is tied up in unused inventory. The books that are on the shelves that are not sold must be shipped back to the publisher, and the publisher reflects that cost in their prices to bookstores. With the Amazon.com model, a single warehouse serving all of the United States—in fact many areas of the world—was able to cut out almost all of that 35 percent.

   Another good example of a solid business model is America Online (AOL). AOL's value proposition is much more subtle. Its business model is to use technology to create an Internet portal with repackaging of information coupled with a simplified user interface that opens up the World Wide Web (WWW) universe to the average consumer.

   Part of the reason that the incomplete business models received funding in the first place is a result of investor frenzy. Many venture capitalists got their money from the first wave of the dot-com craze either as investors or owners. They developed a sense of investment invulnerability, a Midas-touch self-image. For those who had not been involved before, there was a bandwagon effect. The new venture capitalist and the new investors were exercising less due diligence and relying more on instinct, anxious to be able to reproduce the magic of the first round of dot-com successes. Unfortunately, the investors and the new venture proponents began to believe their incestuous string of market projections. And that led to many companies being started that really had no business being in business.

2. *A second reason for collapse was hyperextension.* The hype of unlimited growth and easy money drove companies to invest in expanded infrastructure or corporate acquisitions far beyond the limits of reasonable returns (given today's 20-20 hindsight). Existing carriers or operators, or new carriers aggressively deployed infrastructure using new technology that did have significant cost advantage in the long run. This includes fiber optics, DSL, and fixed wireless technologies. Examples include Global Crossing, NorthPoint, Teligent, and Windstar. The market-cap hype gave them the financial backing to deploy infrastructure at several times more capacity than was reasonable for them to capture considering projected demand and feasible prices.

   Incumbents with existing infrastructure seemed at a disadvantage. They were saying, "How can we compete with these new guys? Our infrastructure is old. Our infrastructure is limited. Our current infrastructure can't compete on a cost to serve basis against these new players." Today they're feeling lucky that they were slow to react.

   The other area of hyperextension consists of companies that wanted to use acquisitions to expand their position in the value chain. Telecom network operators were buying service providers to gain presence in new markets, or network operators or service providers were buying system integrators to get more of that value-chain dollar. They had the money to do it because they had the inflated stock prices. They had the alliance capital. But they lacked the strategic plan and the post-merger integration capability to capture the value from the acquisitions. A good example is PSINet.

3. And last, but not least, *poor execution across several key functions was a major contributing factor to failures.*

   Product Definition, Marketing, and Sales
   I differ with many in the press in terms of why Iridium collapsed as a business. I don't believe that it was because the business model was flawed. I believe it was because execution was inadequate. They knew at least two years before they launched service that they had to have a dual-mode phone capable of penetrating buildings using an existing cellular network, while capable of working outside of cellular coverage via the satellite network. They knew that phone had to have a price of less than $500. Their extensive market research made this clear. But after spending almost $4 billion on the network, their Board refused to spend an incremental $25 million to develop an acceptable handset in time.

   Iridium and its regional investors also failed to execute in marketing and sales. The concept was born in the euphoria of the early 1990s, when cell phones were an oddity, and the view was that people would beat a path to their door to get service. But the service arrived in the market in 1998 when most wealthy people in developing countries had cell phones. The subset of people who needed the additional satellite service coverage had to be actively marketed and sold, but the effort was not there. Teligent and Windstar also had issues in the areas of marketing and sales, and product definition.

   Provisioning and Customer Service
   Provisioning and customer service were also fundamental causes of failure. The DSL CLEC failures in many cases are tied to these types of fundamentals. One of our clients had a 90-day waiting list for service for new customers. They couldn't provision fast enough despite increased personnel because they hadn't really worked out all of the difficulties involved in doing so. This was complicated by lack of cooperation

from the incumbent. The CLEC, however, could have foreseen most of the complexities and been prepared. As the problems arose, the CLEC didn't react well to the complexity and wasn't prepared for process improvement. The problems were fixed with a process redesign, but not before six months of prior marketing and sales efforts were wasted and a host of disgruntled early adopters dampened further sales growth.

There are other issues with customer service. Many service providers implemented new technologies but did not anticipate the needed changes in customer-service approach. By default, the customers were the guinea pigs and the customer-service people were unprepared to handle both the volume and complexity of service issues. Thus customer satisfaction declined and churn rose leading to slower than expected net revenue growth. They could make the sales, but they had to overcome growing churn to achieve net growth. Many of the newer mobile wireless carriers suffered (and continue to suffer) from this problem.

As a result of slower than expected revenue growth, everybody's stock price was down. Investors began to have second thoughts about further investments. There was admitted self-doubt on the part of the smarter major investors as they became more conservative because they had made calls on things that didn't pan out.

Many of the smaller, established companies ended up in Chapter 11 even after significant Draconian cost-reduction measures. Even the larger companies were forced to reduce head ranging from 30,000 in a giant company such as Alcatel to thousands in just about every good-sized telecommunications company, including major carriers such as Sprint.

In many cases, investments were deferred or cancelled. Research and development (R&D) or product development, or pilots or new network rollouts, were reduced in scope or stopped.

To compensate for slower revenue growth, capital-management programs were instituted looking at inventories, receivables, and supplier payment contracts. Suppliers were rethinking those vendor financing agreements that only 24 months prior had seemed like such a great idea in backing the network of a new upstart company.

And last but not least, the companies that were going to survive and felt comfortable recognized that they had a good product and that customers would be willing to pay a little bit more for it.

How to Weather the Storm

In the following there are some guidance to service providers, network operators, and equipment manufacturers on what should be done until the market sorts itself out.

What Should Service Providers Do?

- *Streamline their operations.* Go back to the basics of your business: provisioning. The waiting lists are still too long. Some installers are untrained, inefficient, and inept. There are too many callbacks, which hurt profitability. The service-quality complaint levels are still too high. In many cases, the provisioning processes are not well designed. Get customers provisioned quickly, correctly, and cost-effectively.

  Self-provisioning is a great concept. Many in the industry are working on the issue. It is a terrific idea. But it will need time to be fully developed, such as Microsoft's Plug-and-Play concept. Plug and Play has been around for several years as an idea. But even now adding a new piece of hardware or software to your computer embodies an element of risk. There's a big gap between concept and reality. Likewise, self-provisioning is a wonderful concept, and working toward it will be a rewarding task—but it is not going to happen overnight. And it is not going to happen easily. Keep at it.

- *Customer service should get back to the basics.* Take the time now to clean up customer service and turn it into customer-relationship management (CRM). The data in the industry points to poor customer service as the number one reason behind churn. In churn surveys, people say, "It is my coverage. It is my bill. It is my quality of service." Those are the reasons for the first customer-service call. Frustration builds as, during the first call, it takes several minutes to speak to a human being; then the human being asks you to repeat all the information you just gave to the automated voice response system, (causing increased frustration); and then you are put on hold for several minutes. If the problem is resolved favorably at this point, the service provider probably lives to fight another day. If the problem is not resolved, few have the patience to try again with the same service providers.

  Customers value their own time. Most don't have an assistant to do all this dirty work, so prompt and accurate service response is essential. Voice response systems have saved costs, but some implementations have not resulted in better service. The CRM processes and system requirements should be reviewed against the current metrics of churn, customer satisfaction, and cost.

- *Trim the service portfolio.* Over the years, service providers have developed a broad set of services. Look at what is selling and what is not. Try to understand why and why not. Look at what you are making a profit on and what you will be able to make a profit on. Make the hard calls, because you probably can't afford to support all of those services any longer.

- *Sharpen your customer focus.* Review customer profitability. Think about customer lifetime, customer acquisition costs, the full life cycle, and the total cost and profit of ownership (TCPO). View TCPO from the customer perspective. Think about the total cost and profit of service, and understand the needs of your customers. Take the time now to really understand your customers. Your customer data is a wealth of market information at your fingertips. You are probably not getting the best use of this information today.

- *Get creative on pricing and bundling.* But do it smartly. Bundling is a great concept, but it needs to be done to reflect the customer needs. Bundle to reflect true economies in your pricing and reflect the true value added to a customer of the combined service. The concept of "one bill" is sometimes the only excuse for a bundle. This doesn't cut it alone, not with the business and consumer customers that we have polled. They are much more sophisticated than that. They recognize that in many cases on the Internet, they can create their own bundle from multiple vendors and get automatic billing to their Visa or AMEX card for any combination. So, "one bill" is seamless to them. Let your prices reflect your cost to serve, the higher volume that you can achieve through pricing campaigns, forward pricing, and truly enhanced customer value in the context of the new competitive situation.

What Should Network Operators Do?

- *Replan for network expansion.* If you are operating a network today, your engineers know more about what is going on in the market than the best of the market research companies. They know the hot spots of the growth of demand, and they can tell you how much of that growth is coming from increased use by your current subscribers versus use by new subscribers. These are important vectors. It is the best information you can gather of where things will be going during the next six to 12 months.

  This information can also let you know where are you under-utilized and where are you over-utilized. This information should be used at the strategic planning level, not just for network planning. Network planners or network engineers need to have raised visibility in the company, because they have information that has important strategic corporate value.

- *Implement stop-gap measures.* Where can the infrastructure be upgraded rather than replaced? Where is a phased replacement or upgrade necessary and possible? How can under-utilized elements be reused or redeployed? The engineers can figure out what is possible. Again, engineering can add value at the corporate level.

- *Make the right next move.* Carefully plan your next major deployment. Vendors may be disappointed to delay that big sale. Carriers may have to wait to solve all of your current network's problems with the next-generation network. Choose the move that makes the most sense for your markets and your financial situation.

- *Make the next move right.* Think of this as an opportunity to really plan what you are going to do next. And not just the deployment of a network, but also how you are going to sell it, how you are going to provision it, and how you are going to provide customer service to it. Take a holistic look to provide a superior user experience. Select the right technology and develop a strategy for trials and pilots of the processes and systems. Take the time to ensure the quality of the end-user experience.

What Should Equipment Vendors/Network Integrators Do?

Right size your organization to the new market realities.

- *Trim the product lines.* Can you afford to have all of your current products in service at once? Rethink the R&D pipeline. Think about an old term, P^3I Preplanned Product Improvement, because the carriers are thinking that way. They're thinking about how to get by with lower capital expenditures (CAPEX). You are going to need to right size your capacity, whether it is the factory floor, fabrication facilities, R&D, product engineering, or system integration. You are going to have to right size it because explosive growth has stopped. Prepare for a more modest pace of growth in the industry. Those of you who are the survivors, there's a wave of consolidation about to occur, which will lead to subsequent right sizing until a profitable balance is again established.

- *Optimize capital management.* Look at inventories, supplier relationships, etc. There are a lot of things that can be done to free up capital and reduce operating costs.

- *Review supplier relationships.* How many different suppliers are required? Understand the difference between a vendor and a strategic supplier. A vendor is one of several suppliers capable of providing you a similar commodity at a similar price. A strategic supply partner has unique capabilities key to your business success. You can work together in product R&D and inventory control to support mutual future success. If you hadn't already put these on a strategic footing, now is the time to do that.

- *Nurture customer relationships.* The service-provider and network-operator customers that you have today should be the customers of tomorrow. And although they may be buying only small bits and pieces over the next 12 months, in many cases, those bits and pieces may be the foundation for their next major CAPEX. It is wise to selectively invest to help your customers really understand the value that the features and pricing of your new technologies can have in making those customers succeed with their customers.

Summary

The telecommunications and Internet market opportunities are still there. The fundamentals of the industry are sound. Five years ago, a 10 to 15 percent per annum growth was projected for service providers for this time period. Although these estimates were viewed as conservative or pessimistic two years ago, they are probably the baseline today. The problem is that during the last two or three years, people geared up for 20 percent per annum growth, which I don't think is going to happen.

The technology and business capabilities are just as great, but they have value-chain implications to be understood. Technologists, engineers, and product developers are developing new products that still obey Moore's Law. That is, they can create longer-lasting products with more capacity, more functionality, and more capability at lower cost. So, if the service-provider revenues are going up at 10 to 15 percent per annum, network operators will likely be in the 8 to 12 percent range, and equipment manufacturers will be in the 5 to 8 percent range. Those with better service concepts, better products, or better execution will rise above the averages as they gain market-share.

During the first quarter of 2002, there was a transition from a stop to the more modest growth rates just noted. In this lull, the industry has time to do some of the things that were neglected in the rush of the past three years. Regroup. Take stock of where you are. Rethink your strategy, where you want to go, and how you want to get there. It is time to establish a foundation for future growth.

Now, which companies will be the future winners? As for specific companies, I don't know. Hopefully, some of the recommendations that I have given will be useful to achieve success. The best guess is that the winners will be those that do the following:

- Conserve their current resources
- Develop an honest assessment of where they are
- Develop a clear strategy built on their unique market position
- Create a practical implementation plan to execute the strategy
- Persist in execution while learning and adapting to change

Scarcity and Complexity

Re-Defining Carrier Networks in the Age of Bandwidth and End-to-End Ethernet Solutions

Philip K. Edholm

Chief Technology Officer and Vice President of Network Architecture, Enterprise and Optical Ethernet

Nortel Networks

Executive Overview

In any system, there is a balance of complexity versus scarcity. When things are scarce, we implement complexity to allocate that scarcity to those elements requiring it the most. When things are plentiful, the allocation systems can be simple. One of the most interesting changes in any environment is when something that is perceived to be scarce becomes plentiful (or vice versa), upsetting perceptions about allocations and complexity. Today we are experiencing such a transition—from bandwidth as a scarce resource to plentiful capacity

The evolution of computing from mainframes with card decks to mini-computers to desktop personal computers (PCs) is an example of a transition from scarcity and complexity to distribution-reduced controls over a 30-year period. One of my professors, when I was a young computer-science student, related the following: When building 10 units of a system, you buy all the hardware you can afford to minimize the investment in software development; when building a million units, you optimize the software to minimize the cost of the hardware. While this concept is obvious to the most casual observer, the challenge is what to do when you are building 10,000 units—obviously, we are trading complexity in one part of the system to minimize scarcity or cost in another. Ongoing assumptions about relative scarcity and complexity works well in a world that is changing slowly, but these assumptions need to be re-examined when scarcity changes dramatically (the advent of application-specific integrated circuits [ASICs] and large-scale field programmable gate arrays [FPGAs] reduced the custom hardware volume point dramatically).

This paper anticipates that fundamental changes in network elements and traffic will dramatically shift the balance of bandwidth, complex control, and management that currently exists. The advent of highly variable data-rate edge traffic of multiple classes over a consistent end-to-end Ethernet/Internet protocol (IP) structure and the explosion of cost-effective bandwidth will enable a fundamental redefinition of the relationships of complex and scarce elements. The result will be a dramatic transformation not only in how we invest in network equipment, but also in how we manage and provision those systems. While this transformation has already begun in the Internet world, it will have dramatic impact on the more traditional enterprise networks as well as on the overall network infrastructure. This paper investigates these trends and their impacts on network architectures, provisioning, and management.

The Bandwidth Phenomenon

Traditionally, the mechanism of delivering services over carrier networks has been to implement circuits that provide connections between computers and sites. (In fact, the applications were designed to run on point-to-point circuits.) As data communications has become more prominent, new technologies, such as frame relay and asynchronous transfer mode (ATM), have emerged that are more suited to the bursty nature of data applications, while maintaining the circuit topology. However, at the core of today's network philosophy are concepts that tie together edge services, circuit provisioning, and core engineering into a tightly interwoven set of functions.

Within these networks, the management and operation of the network are defined based on a combination of resource availability and traffic characteristics. In the voice world, all circuits were of relatively fixed bandwidth, and networks were designed to accommodate that need. As data services became more prominent, the traffic variability led to new services, such as frame relay and ATM, with variable rate capabilities. At the same time, bandwidth continued to be a critical scarce resource in networks. This was made painfully clear as the demand laws of computing began to drive the network service demands. While voice demands are driven by simple growth in the number of users and average use (as a voice channel is "fixed" at 64 kilobits per second [kbps]), data demands are driven by Moore's Law (silicon doubles in capacity every 18 months—and applications and files follow suit), Metcalfe's

Law (the value of a network is an exponent of the number of users, and the traffic grows at an exponential rate—witness the number of Cc:'s on your recent e-mails), and major transformational events (client-server computing, the Internet, etc.). The result is that bandwidth demands since 1990 have been increasing at two to three times per year (and even more in the explosive time of Internet growth). The result has been a scarcity of bandwidth at many levels of the carrier infrastructure.

Coming from a circuit perspective and a scarcity of bandwidth leads to a plan of network operation that optimizes network implementation and operation to accommodate those parameters. So most carrier service networks today provision individual circuits and use that information to mange and allocate end-to-end bandwidth through the network. The concept of "end to end" results in a complex management task at both the edge of the network and within the core elements In other words, we trade complexity of operation and control to provide a guaranteed service and end-to-end provisioning of that service. The result is a minimization of scarce resources (bandwidth and associated components). This philosophy has extended to new networks as well. Both ATM and IP networks have employed a mechanism to reserve bandwidth on an end-to-end basis (resource reservation protocol [RSVP] for IP). The assumption for this type of mechanism is that bandwidth is scarce and must be optimally provisioned and allocated and that the average traffic requires service-level agreement (SLA) parameters guaranteeing delivery and latency.

An interesting side point is that the Internet, emerging as it did from an any-to-any environment, was operated and managed (within the fashion that the technology enabled) as a network where the core infrastructure was provisioned and controlled separately from the edge service connections. However, as the network was designed as a best-effort system, the capability to deliver more sophisticated service guarantees did not exist.

It would appear that we are on the verge of a profound transformation in how the networks are operated and provisioned. This transformation is based on three convergent forces: plentiful bandwidth, convergence to IP/Ethernet traffic, and multiple service classes that will dramatically change the way in which we think of network structures.

Through recent technological advances, bandwidth is becoming plentiful (while not free). This is the result of the explosion in optical bandwidth brought about by dense wavelength division multiplexing (DWDM) and new laser technology (jumping to 40 gigabits per second [Gbps]) as well as fundamental changes in the requirements and devices required for forwarding at network connection points. The convergence of traffic onto Ethernet (99 percent of all enterprise traffic starts and ends as Ethernet) and IP (now the dominant protocol) has reduced the complexity of the devices. As the tasks required have become stable, they can be moved into silicon. In fact, the advent of multiprotocol label switching (MPLS) (or generalized MPLS [GMPLS]) as a forwarding mechanism in the transport core holds the promise of delivering an exponential increase in bandwidth over the emerging optical capabilities. Optical bandwidth has been doubling every nine months since 1997, and the cost per bit has dropped by 50 percent every nine months in the same period. Together, this represents an increase that is eight times greater than Moore's Law and is increasing at about six times per year. The movement (begun in the enterprise) of using highly specialized ASICs and network processors to manage the traffic at intersection points appears to be capable of keeping up with this demand. The result is that we have a "bandwidth glut" for the first time since the advent of the computing age. This trend is extending to the edge of the network with the recent explosion of optical Ethernet in metro networks. By using Ethernet both as a basis of service offer and as a transport element, the bandwidth available in the last mile and in the metro is increasing exponentially. Service offers of 100 Mbps or 1 Gbps are now being delivered. When combined with the existing enterprise campus bandwidth (typically 100 Mbps per desktop), this portends a dramatic change toward bandwidth abundance.

The Traffic Phenomenon

The explosion of bandwidth demands, as driven by computing engines, has very different characteristics than previous traffic. It is highly bursty and cyclical. The variability at the edge of the network is extreme. Also, to support the next generation of applications (such as streaming, storage), bandwidth operational rates are much higher, further exacerbating the problem. At the same time, applications are converging to common transports (primarily IP and Ethernet). Where previously separate networks were often deployed for voice, mainframe data, e-mail, etc., the convergence brings all of these traffic types into a single network interface. This enables a re-thinking of the trade-offs between service types and how the traffic is managed across the network.

As different traffic types converge, the variability of requirement can be used to manage their individual demands, simplifying the network. A typical SLA can be defined with four variables: bandwidth, delivery guarantees, latency guarantees, and availability. By varying these four parameters, different classes of traffic can be easily defined that meet different applications needs. For example, voice requires relatively low per-user data rates but has high requirements for delivery and latency guarantees as well as high availability requirements. A contrasting best-effort service, potentially used for Web browsing and non–real-time apps, such as e-mail, would have higher bandwidth needs (especially instantaneous), while being able to use much lower delivery and latency guarantees and lower availability. Mission-critical data lies somewhere in between.

Convergence and Change

The impact of the three convergence elements (bandwidth, bursty, and different traffic types) will be dramatic. *Figure 1* shows a simplified representation of a projection of traffic types from 1980 to 2010 using percentage of total traffic. The vertical axis is percentage of total traffic. As can be seen, in 1980, 98 percent of all traffic was voice, and, by 2010, less than 10 percent of traffic will be voice (including videoconferencing). In addition to the voice/data transition, there is another factor that is impacting the traffic model. Prior to 1995, virtually all data traffic was mission critical. In other words, it was traffic that was absolutely required to run a business or other function. However, in 1995, the advent of the Web browser and the World Wide Web (WWW) changed that paradigm. Casual Web browsing and attendant bandwidth uses have created a new class of traffic, commonly known as "best effort." *Figure 2* anticipates the growth of best-effort traffic impact on the model. As can be seen, now the traffic percentages are not only becoming predominately data, but the best-effort traffic is also becoming a significant if not dominant part of the overall traffic. Finally, prior to the bandwidth explosion, networks were operated to maximize bandwidth fill. Significant activities were undertaken to assure that the network was operating at 80 to 90 percent of capacity whenever possible, maximizing the return on scarce resource investments. However, starting in about 2000, the impact of the bandwidth explosion has been changing the fill requirements. In fact, as the cost structures of DWDM, 10-Gigabit Ethernet, and next-generation MPLS systems impact, the cost of increased bandwidth becomes much less. The result is shown in *Figure 3*. Prior to 2000, the network capacity was only slightly above the demand, but that gap is increasing (or will increase) where the new technologies, such as optical Ethernet, are available.

The result of these changes is a rethinking of how services are delivered and provisioned, and of how the network is managed. Systems such as ATM and RSVP (see *Figure 4*) were defined when the relative balance of traffic types and fill required extensive management and reservations to deliver cost-effective solutions. However, as we look forward, we see a dramatically different world. As the different traffic types each have different requirements, by properly classifying the traffic and then managing the operational parameters of the infrastructure, we can easily deliver a SLA that meets the customer requirements. For example, by giving the small percentage of traffic that is voice a traffic class that has absolute priority of transmission and no discard, the class will be transmitted through the network without any impact of the volume of traffic in other classes. As can be seen from Figure 4, the combination of the emergence of, first, mission-critical data and, next, best-effort data, when combined with the cost-effectiveness of operating the network at a lower fill percentage will make the management of diverse services easy. This is further emphasized by looking at the trend over the past five years in enterprise campus local-area networks (LANs), where exactly these mechanisms for classification and management using simple class of service (CoS) without guarantees are common. The failure of ATM and RSVP to build campus LANs is a result of the ease of managing multiple service types (including real-time) over a network that has over-provisioned bandwidth headroom and simple CoS classification and queue-management operation. As shown on *Figure 4*, the advent of simple differentiated services (DiffServ) in the Internet Engineering Task Force (IETF) anticipated these changes.

FIGURE 1

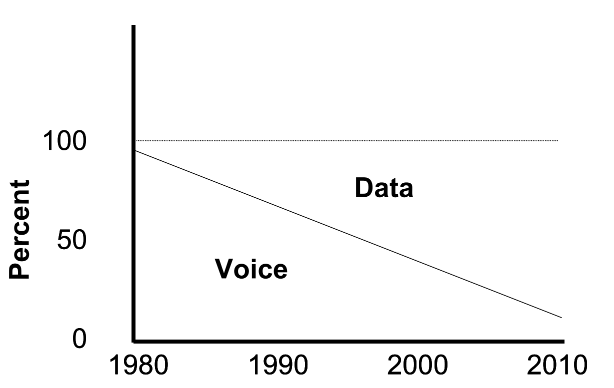

FIGURE 2

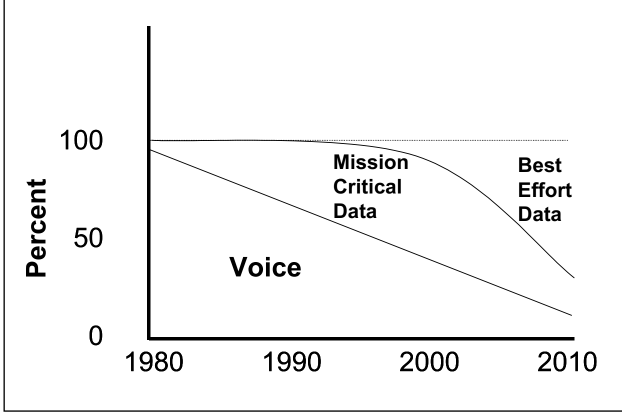

FIGURE 3

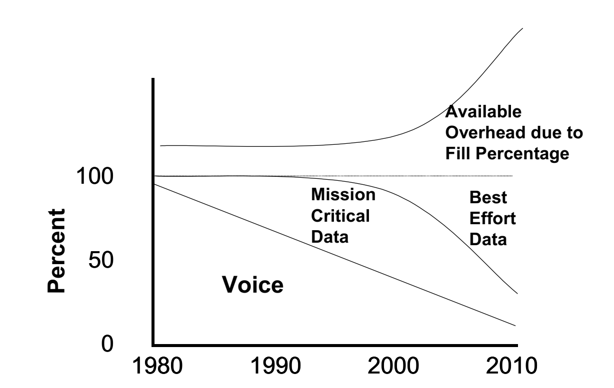

FIGURE 4

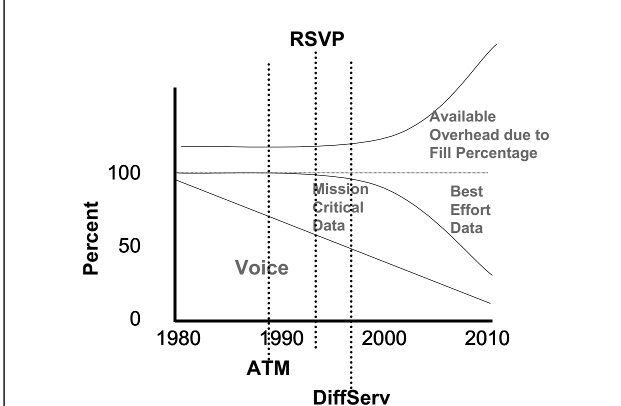

Reducing Complexity

By taking advantage of the changes discussed earlier, it may no longer be necessary to manage the network as end-to-end connections. In fact, managing the network as an access edge and a separate core will be a major change that will be enabled. *Figure 5* shows the demarcation of the customer (CE), the provider edge (PE), and the transport core. At the edge of the carrier network, there must be a user network interface (UNI) that defines the service characteristics on the interface to the customer. The UNI defines the forwarding mechanism (Layer 2 or Layer 3) and the SLA characteristics (bandwidth, delivery, and latency guarantees and availability). The SLA must be managed and provisioned on an individual port/customer basis. However, once the traffic has left the UNI and entered the transport, using the aforementioned techniques, traffic can be managed using simple techniques. As shown in *Figure 6*, this creates two distinctive management systems. The edge system configures and manages the UNIs based on the services purchased and the associated SLAs. Traffic is classified according to the SLA and input to the transport core. Within the transport core, individual flows are not the critical element. Rather, the aggregate rates in any given class on any given link become the key criteria. Therefore, while the edge is managed as individual customers, the core is managed based on link operation and predictive intelligence that will anticipate resource utilization. This may include predictive systems that can predict the impact of a new customer entering the network based on the locations and SLA parameters. This will enable checks of operational impact prior to allowing the new customer on the network (or allowing request-to-change SLA parameters).

The result of this stratification is dramatic reductions in operational expenditure (OPEX) while having minimal capital expenditure (CAPEX) impact. Today's carriers spend about 60 to 70 percent of their expense in the OPEX category. The explosion of bandwidth, combined with the complexity of multiple traffic types on the same interfaces/networks, will further increase the complexity. If today's mechanisms of individual provisioning and management across the entire network with end-to-end operation for each circuit (or virtual circuit) continues, OPEX will only grow, both in absolute dollars and as a percentage. However, if we assume that bandwidth costs are becoming nonlinear (doubling the bandwidth does not double the cost as it is the same fiber, the same chassis, etc.), a reasonable CAPEX will change this. In other words, spending 25 percent more in CAPEX will generate a 25 to 50 percent reduction in OPEX by stratifying the UNI service edge from the core. The results are as follows in *Table 1*:

The result is a 10 percent reduction in relative cost, while delivering services for which enterprises have indicated that they will pay incremental pricing (10–20x the bandwidth at 1.3–1.5x the price). In addition, the reduction in CAPEX to purchase complex end-to-end provisioning and manage-

FIGURE 5

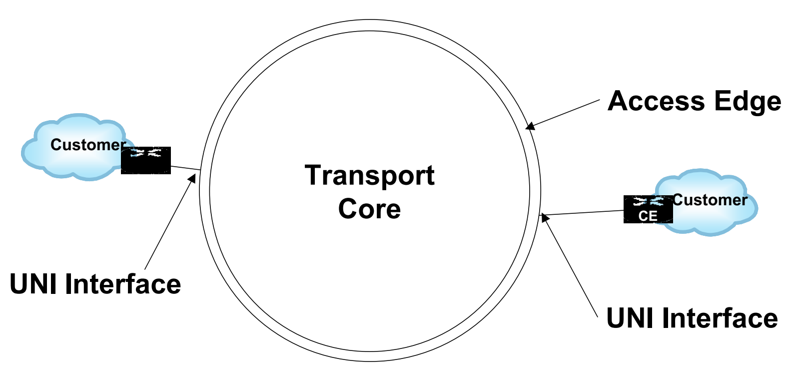

FIGURE 6

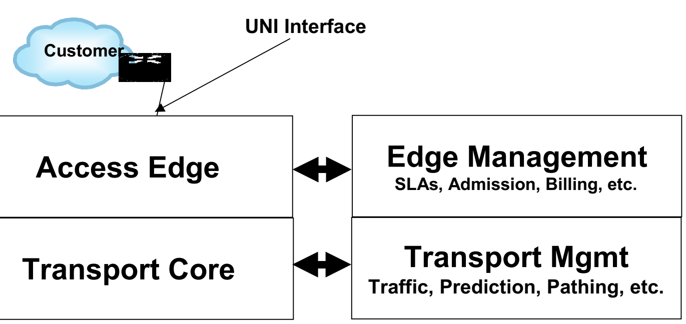

TABLE 1

	Today		*Tomorrow*
CAPEX =	35%	x 1.25 =	43.75%
OPEX =	65%	x 0.7 =	45.5%
Total	100%		89.25%

ment systems may in fact equal the added cost for transformation bandwidth, delivering another 5 to 10 percent of expense reduction. According to a study done by McKinsey & Company and Goldman Sachs, given the rapid growth of data traffic demand over the next few years, unless service providers are able to reduce OPEX by 25 to 30 percent per bit carried for each year through 2005, no reasonable amount of CAPEX reductions will allow them to meet targeted return on invested capital.

Architecting the Future

Delivering the aforementioned capabilities and benefits requires thoughtful analysis when defining architectures. Within the context of the network, the UNI is a thin layer at the edge, only used to define and manage services and assure that SLAs and classifications are properly managed and policed. The rest of the network can be thought of as the transport. Within the transport infrastructure there are forces that drive certain allocations of resources and intelligence. From a CAPEX perspective, *Figure 7* shows cost allocations based on cost per port/customer and per unit. On the left, costs allocated by port/customer are primarily at the edge. This is because each customer must have separate port, access, fiber, etc. However, these edge devices are individually inexpensive. As we move further into the network, individual devices become much more expensive, but they are far fewer in number and amortized over all of the network customers. At the very core, long-haul fiber is incredibly expensive but is the amortized base for all services. To minimize the cost of an infrastructure, it makes sense that (as the arrow shows) intelligence and complexity should be pushed up in the network. This "centralization" reduces complexity, CAPEX cost for the devices at the edge, and amortizes the function over the user base.

Figure 8 shows the same model but for network bandwidth. Individual bandwidth is greatest at the edge, especially in data networks where peak to valley edges can have 3–10x variations. As we move back in the network, the smoothing effect of statistical variations causes the per-customer bandwidth that must be supported to decrease. However, as shown on the right, the link speeds are highest at the core (and the attendant packet processing rates). For example, a network of 100,000 100 Mbps user ports (10 petabits), with an average SLA of 20 Mbps and a average use of 5 Mbps, would need a core capacity of only about 1–2 petabits. Assuming some duplication for capacity in failure and some inefficiencies of the core, the 2 petabit core is reasonable. Assuming this was implemented using 10-gigabit links, there would be 200 links in the core.

While cost tends to push complexity and intelligence up form the edge, achieving the bandwidth required tends to push it down from the core. Complexity is the bane of achieving speed, and complex decisions are virtually impossible at very high speed. The result is that virtually all networks are architected as shown in *Figure 9*. There are three layers: an inexpensive simple edge, a fast simple core, and a layer of intelligence between. This architecture often is depicted as a set of concentric circles, as shown in *Figure 10*, but the result is the same. This partitioning of functionality enables the operation of the network into a logical arrangement that optimizes both cost and performance. Past networks have either been designed or have generally evolved to this structure over time. For example, a frame-relay network has simple edges (point-to-point circuits—not necessarily low cost due to tariffs) connected to an intelligent node that provides the service structure. These intelligent nodes are connected together through ATM or synchronous optical network (SONET) structures that provide fast unintelligent transport.

Building the next generation of networks will involve rethinking our assumptions, as previously discussed, but doing it within a framework of an allocation of resources.

FIGURE 7

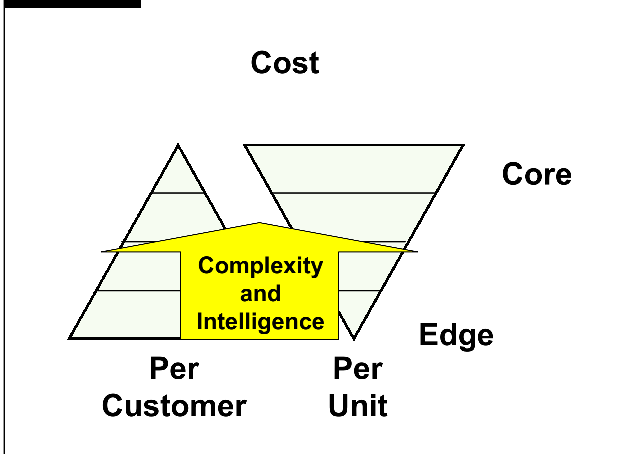

FIGURE 8

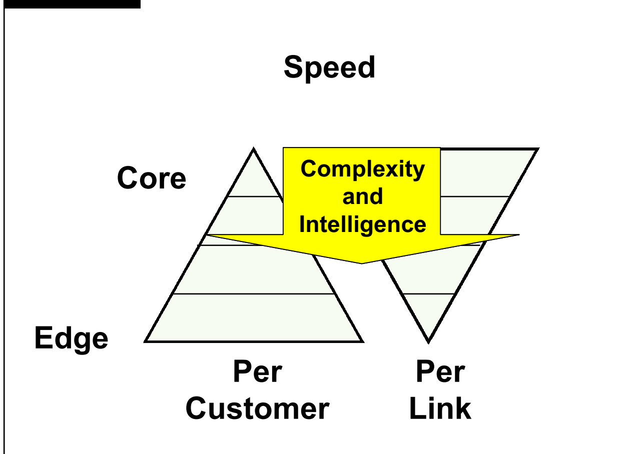

FIGURE 9

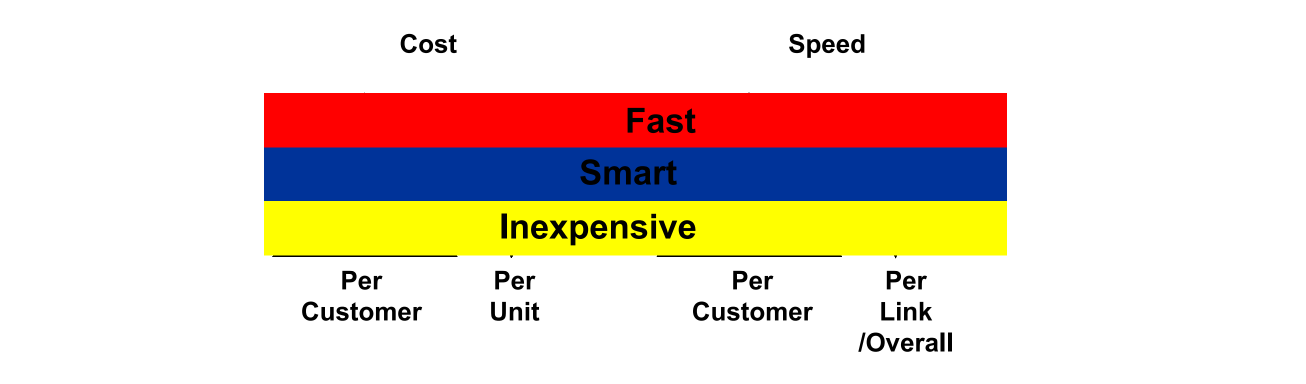

One critical element of this thinking is the relationship between the metro and the long haul. The emergence of MPLS and GMPLS along with high-speed optical appears to define the core of these next-generation networks. The challenge is how to intelligently connect these cores to the customers. Optical Ethernet, delivering services (either Layer 2 or Layer 3) over Ethernet connections, seems to be the preferred mechanism for next-generation edges. This solution allows the enterprise traffic (which is 98 to 99 percent Ethernet in the campus today) to remain Ethernet end to end. While traffic will remain Ethernet end to end, decisions in the processing of packets will be made a multiple layers. *Figure 11* shows the structure of the lower three layers of the network model circa 1990. As can be seen, the early explosion of creativity and invention in networks created a large number of competitive (or at least co-existent) protocols and media/control.

During the late 1990s, organizations converged on IP as the common Layer-3 protocol, and within enterprise campuses Ethernet became the dominant solution. *Figure 12* shows the 2001 protocol map. In all technology solution sets, each layer tends to converge on a de jure or de facto standard. The same convergence seems to be happening at Layer 2, where Ethernet is fast becoming the converged standard. This is based on the enterprise preference, the end-to-end values of commonality, and the cost of the components. As this convergence happens, the future protocol will be as shown in *Figure 13*. This convergence has significant impact on the way in which we process and move data. In the highly divergent networks of the past, where the layers implemented different standards, systems had to be built to do functions on each layer separately. As we consolidate, decisions can be made simultaneously between the layers. For example, a forwarding decision can be made at Layer 2, while a quality of service (QoS) function can be at Layer 3 using DiffServ. This blurring of the layers dramatically simplifies networks and will lead to the next generation of products—the specialized system. When the complexity in networks was matching and managing complex dissimilar protocols, the key was devices that dealt with the maximum in protocol and interface diversity. As we move forward, specialized devices, optimized for a specific point in the network and functional requirements limited to functions in Ethernet and IP at that point, will dominate. This has happened in the enterprise campus, and it is happening in virtual private network (VPN) nodes, in Web switching, etc.

Within this framework, the challenge is to allocate the intelligence to deliver services in a way that optimizes both cost and speed. Cost in this context is both acquisition cost (CAPEX) and ongoing operational cost (OPEX). As we have discussed before, trading bandwidth for complexity can generate significant returns in reducing operational costs. With these concepts in mind, an architecture for the future can be postulated with a distinct separation between the metro transport and the long-haul transport. Within the metro, an Ethernet transport yields significant benefits as it

FIGURE 10

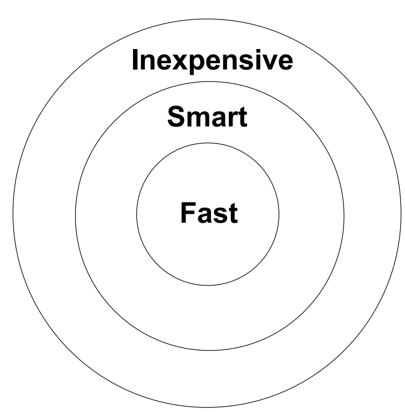

FIGURE 11

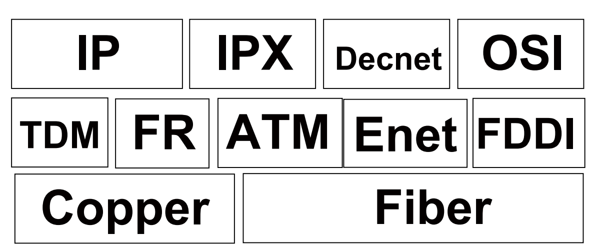

FIGURE 12

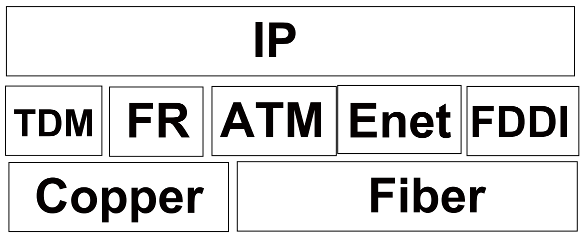

FIGURE 13

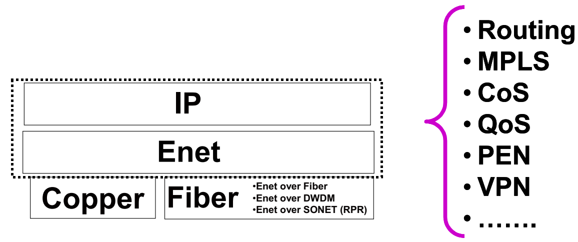

enables bandwidth to minimize the complexity of path and QoS decisions at the low-cost points generated from the enterprise LAN segment. Beyond the metro, into the long-haul environment, the simplicity of the Ethernet transport mechanism does not offer the necessary benefits of scale and traffic management. Therefore an architecture with the elements as shown in *Figure 14* combines the best of all worlds. In this architecture, the metro edge becomes a distributed or "logical" provider edge connected with a Layer-2 Ethernet transport to nodes that interface between the simple Ethernet metro transport and a highly scalable, traffic engineered MPLS core. The three elements of the logical PE are the LPE edge, the LPE Layer-2 switched transport, and the LPE core. In this architecture, functions can be allocated based on the previously discussed requirements. The LPE edge device has responsibility for SLA management and customer-service generation. The Layer-2 transport is over-provisioned for simplicity; the LPE core manages the complexity of the MPLS core network.

Figure 15 shows how the earlier layering of low-cost, smart, and fast can be applied to the resulting infrastructure. As can be seen, this solution optimizes all of the necessary elements. Service management is at the edge, where it belongs, managing the actual performance of the service delivery and enabling SLAs that can deal with the incredibly bursty nature of data. The metro transport is optimized for speed without being burdened with the complexity and intelligence necessary to provide large-scale national networks. The smart layer at the interface between the metro and long haul translates from the simplicity and cost-effectiveness of the metro to the more complex and scalable core. Here is the intelligence necessary to deal with the potential thousands of paths and connections necessary to build such networks. Finally, the long-haul core is very fast, based on a combination of MPLS and optical. This provides the high-speed capabilities to minimize complexity at reasonable cost.

The results of this architecture are significant. In the metro, it is possible to deploy the low-cost devices and inexpensive transports derived from the millions of Ethernet enterprise ports and devices. The edge devices can do the policing and management for the SLAs, including assuring customer separation through simple encapsulation and auto-discovery of

FIGURE 14

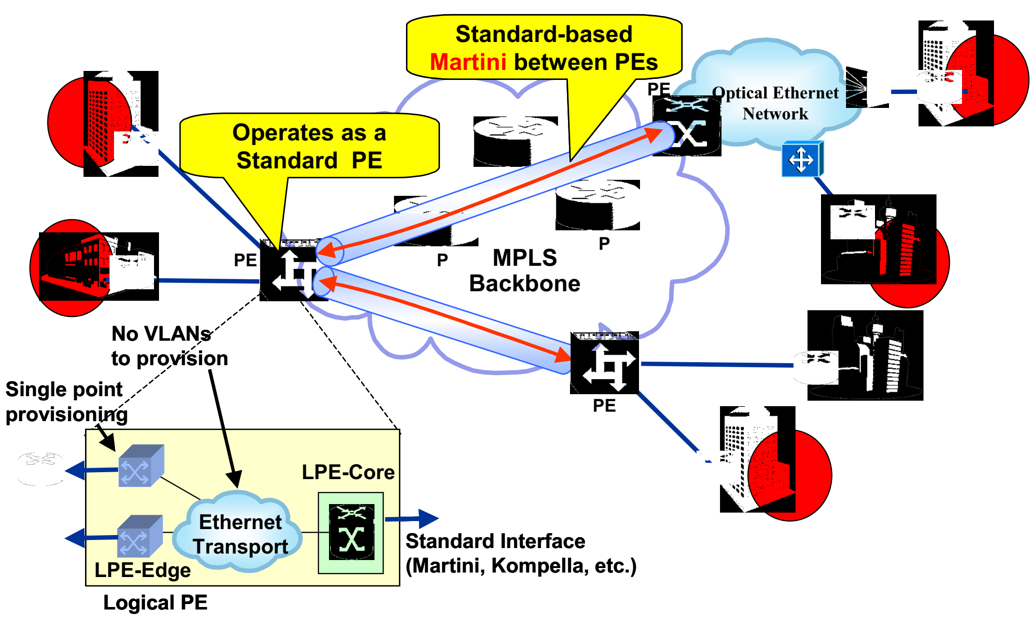

FIGURE 15

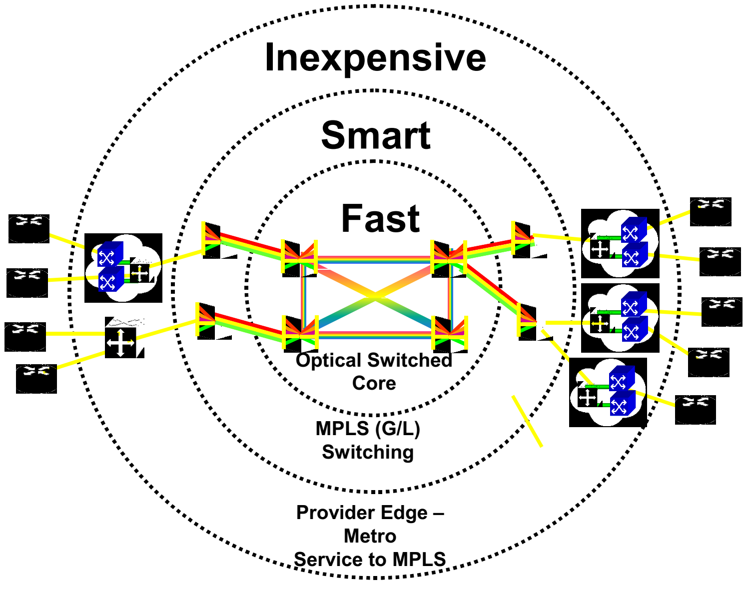

paths to other nodes. Between the metro and the long haul, the traffic is mapped into MPLS. MPLS is the logical alternative, as it enables traffic engineering and management across the core for these new services, while also enabling existing services to coexist through GMPLS. A large number of carriers are committing to MPLS as the way to manage their core backbones. (As the traffic is already separated by customer, the MPLS label-switched paths [LSPs] need only be between the metros, as they can be shared between multiple customers.) *Figure 16* shows the impact of this change. Instead of having to build LSPs for each customer connection or LSPs between all the edge devices, LSPs are only required between the metro (logical PE core) devices. As the number of paths and complexity expand exponentially with the

FIGURE 16

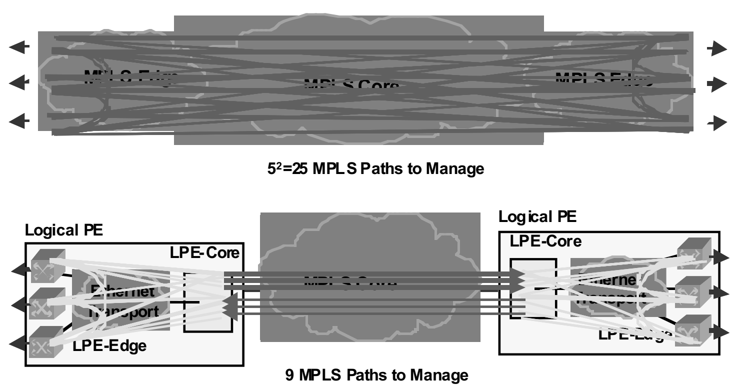

Fan Out	Number of Paths	Complexity
1:1	10n	$100n^2$
2:1	5n	$25n^2$
10:1	1n	$1n^2$

number of points (10 points connected is 10x9 or 90 paths; 100 points is 99x100 or 9,900, or 100 times as many), moving this point to the metro/long-haul edge and having 20–50 metro edge devices connected reduces the complexity by 100–10,000 times.

As bandwidth throughout the network is now much higher and the traffic types are varied, the network can be managed as an edge admission network and a core transport. As was discussed earlier, the edge can be managed for SLAs and customer requirements; the core links can be sized and managed for presented traffic; and the different traffic types will manage any load variation on a real-time basis. The result is a network that is very easy to manage and provision, both at the customer level and the overall transport level.

Conclusions

While the transformations anticipated in this paper will not happen overnight, they are in fact beginning throughout our networks. On enterprise campuses, the bandwidth explosion has already transformed the networks, and DiffServ-based CoS is used for IP telephony. Optical Ethernet, being delivered in some markets and under evaluation by virtually every carrier, will extend this bandwidth into the metro. Finally, the explosion of capacity in DWDM and the emergence of MPLS/GMPLS will deliver the core long haul.

Based on this combination of factors, it is entirely possible to anticipate that the beginning of the new millennium will be viewed in the future as the beginning of the transition from telecommunications networks ruled by the laws of the circuit and the age of scarcity to the new networks ruled by the laws of packets and plentiful resources. Surely the carriers and service providers that anticipate these changes will profit from the inevitable value that they create in the marketplace and the business capabilities that they deliver. The critical element is to understand how to migrate to the new future. Much as the buggy-whip manufacturer who realized his business was selling vehicle control systems prospered by moving to building steering wheels (while his competitors pushed blindly forward trying to sell whips to car owners and failed), the successful carrier/service provider of 2010 will have adopted these new paradigms, realizing that the business is not delivering circuits, but rather effective communications systems that extend the local communications into metro, regional, national, and global domains. A new world awaits us to bring it to fruition, and the transition is as profound as that from analog to digital or the wireless and Internet revolutions. This change will, in the end, redefine our perceptions of time and distance and enable a new set of applications and services that will continue the information revolution.

Telecom Markets and the Recession: An Imperfect Storm

Dean Eyers

Group Vice President and Worldwide Director, Telecommunications Group

Gartner Dataquest

Abstract

In difficult economic conditions, the global market for telecommunications services continues to grow, while the outlook for the equipment sector is mixed—some sectors will remain depressed through 2003.

Predictions

Gartner Dataquest makes the following predictions:

- The global telecommunications equipment market is expected to decline by 4.3 percent in 2001 but will recover and grow by 6.0 percent in 2002.
- Despite negative sentiment in the telecommunications market, revenue from global telecommunications services will grow by about 13.0 percent in 2001.
- The mobile handset sector may be experiencing its first real downturn, but infrastructure investment—mainly in third generation (3G)—is still increasing.

Overview

The current blind panic in the telecommunications market is no more justified than the blind faith that preceded it. Although many sectors and vendors are undoubtedly suffering, there are also bright spots in the gloom—not least of which is the fact that the telecommunications services market continues to grow strongly. The terrorist attacks in the United States on September 11, 2001, have been one factor in the downturn in the telecommunications market. However, in the medium to long term, Gartner Dataquest expects the demand for more bandwidth and for more secure, more redundant, and more distributed networks and services to remain strong. Continuing unstable economic conditions will put many of the weaker players at risk, but other telecommunications companies will emerge in a stronger competitive position.

Markets in Perspective

The telecommunications services market grew in all regions during 2001 and will continue to do so through 2005. Although the market for fixed, circuit-switch voice services is fairly mature, strong growth in the data services, packet voice, and mobile services sectors will combine to drive overall growth to more than 10 percent per annum. The telecommunications services market is generally less affected by economic cycles than other sectors, and there is much evidence that the impact of both the economic downturn and the events of September 11 will be minimal.[1]

Moreover, recent analysis of operator data and end-user plans for telecommunications spending indicate that usage levels and spending will continue to grow even in the regions most affected by the economic downturn.[2] While some of these companies and analysts may have planned for exponential return and radical revenue shifts, the pattern of steady growth is in line with previous Gartner Dataquest forecasts for 2001 and beyond.[3]

In contrast, the telecommunications equipment market has undoubtedly declined during 2001. North America and Western Europe have been hit particularly hard. The Asia-Pacific (and especially China) is the only real bright spot. The fixed infrastructure, enterprise communications, and mobile terminal sectors have all been adversely affected. However, there are still a number of subsectors that are experiencing reasonable growth—for instance, mobile infrastructure, optical metropolitan-area networks (MANs), and interactive voice response (IVR) systems.

The infrastructure segment will remain relatively depressed through 2002 and will pick up only slightly in 2003 as operators absorb the “overbuild” of 2000 and review their plans for next-generation equipment. Gartner Dataquest predicts that there will be a return to strong growth in 2004, as these revised plans are implemented.[4]

The mobile infrastructure market will continue to grow strongly in 2002 and 2003 as operators invest in 3G systems but will grow more gradually after that. The mobile terminal market will “bounce back” somewhat in 2002, buoyed by demand in the Asia-Pacific and by Europe’s eagerness for new products. From 2003, revenue growth in these markets is expected to be more conservative and will be based partly on a shift to higher-end products.[5] The enterprise market will return to overall growth in 2002, as users finally make delayed system upgrades and invest in technologies that will help to support their business against a background of staff reductions.[6]

Taken together, these equipment forecasts represent a large reduction in the expectations of a year ago. A combination of factors—subsequently outlined—have combined to greatly reduce the market potential in 2001 through 2003. However, by 2005, most sectors will return to levels that were forecast previously by Gartner Dataquest.

These forecasts do not support the widespread perception that the telecommunications market in its entirety—or even the equipment sector—is in complete meltdown. Certainly, there are highly visible companies and highly visible sectors that enjoyed very rapid growth during 2000 and are now experiencing dramatic decline. However, looking across the telecommunications market as a whole, the picture is somewhat different. For example, Gartner Dataquest estimates that the telecommunications equipment revenue of eight leading telecommunications vendors declined 11 percent year-on-year for the first nine months of 2001 against 2000. But, a broad-based sample, including smaller Japanese and Asia-Pacific vendors that are closer to the growth markets, shows a year-on-year increase of more than 5 percent. The fourth quarter of 2001 will be a telling one for many companies.

Table 1 shows worldwide telecommunications market revenue by region and sector from 1999 to 2005. *Table 2* shows worldwide telecommunications equipment market revenue by region from 1999 to 2005, while *Table 3* shows worldwide telecommunications services market revenue by region from 1999 to 2005.

Drivers of Economic Depression

Some of the origins of the downturn in the telecommunications sector predate today's general economic malaise. Like much of the information technology (IT) industry, telecommunications operators and equipment vendors were caught up in the dot-com boom. The growth of Internet commerce, services, and revenue has clear and direct implications for those providing the bandwidth and underlying networks.

When the dot-com implosion occurred, it hit the telecommunications industry in several ways, including the following:

- Markets had to radically readjust to the fact that Internet commerce would not grow at 400 percent per annum. Such large growth was never predicted by Gartner Dataquest.
- Stock prices were adjusted downward. The fact that a company is closely associated with the dot-com sector certainly no longer increases its stock valuation; indeed, some would argue quite the opposite.
- Venture capital was severely curtailed and continued only on a very different basis—moving quickly toward profitability is now the order of the day. Those telecommunications services start-ups that still exist independently have generally been forced to reduce spending dramatically and focus on short-term return on investment (ROI).
- Having invested heavily in network upgrades through 2000, incumbent operators, no longer looking over their shoulders in fear of competition from rapidly growing start-ups, could improve short-term profitability and

TABLE 1

Worldwide Telecommunications Market Revenue by Region and Sector, 1999–2005 (Millions of U.S. Dollars)

	1999	2000	2001	2002	2003	2004	2005	CAGR 2000-2005
Asia/Pacific and Japan	$257,384	$307,272	$344,879	$378,146	$410,609	$446,070	$489,163	9.7%
Central and Eastern Europe	$37,499	$42,666	$49,292	$56,879	$63,883	$67,861	$71,479	10.9%
Latin America	$99,429	$116,097	$129,724	$143,220	$161,641	$184,680	$201,667	11.7%
Middle East and Africa	$51,283	$57,646	$64,300	$72,538	$81,659	$93,256	$108,869	13.6%
North America	$377,404	$432,844	$456,763	$497,401	$545,000	$596,490	$650,623	8.5%
Western Europe	$292,805	$316,647	$328,956	$357,231	$378,843	$401,074	$426,304	6.1%
Total Telecom Equipment	$311,302	$381,030	$364,777	$386,482	$410,650	$443,920	$491,770	5.2%
Total Telecom Services	$804,502	$892,142	$1,009,137	$1,118,933	$1,230,985	$1,345,512	$1,456,336	10.3%
Total Telecom Market	**$1,115,804**	**$1,273,172**	**$1,373,914**	**$1,505,415**	**$1,641,635**	**$1,789,432**	**$1,948,106**	**8.9%**

Source: Gartner Dataquest (November 2001)

TABLE 2

Worldwide Telecommunications Equipment Market Revenue by Region, 1999–2005 (Millions of U.S. Dollars)

	1999	2000	2001	2002	2003	2004	2005	CAGR 2000-2005
Asia/Pacific	$66,376	$90,045	$98,511	$111,274	$122,238	$134,532	$150,357	10.8%
Central and Eastern Europe	$13,115	$14,221	$15,048	$17,530	$20,384	$19,983	$20,195	7.3%
Latin America	$20,894	$25,525	$24,017	$24,146	$28,160	$32,207	$33,714	5.7%
Middle East and Africa	$12,632	$14,589	$15,437	$16,721	$16,855	$17,377	$18,345	4.7%
North America	$109,995	$130,428	$118,398	$117,940	$122,531	$135,402	$156,032	3.6%
Western Europe	$88,290	$106,221	$93,368	$98,871	$100,483	$104,418	$113,126	1.3%
Worldwide	**$311,302**	**$381,030**	**$364,777**	**$386,482**	**$410,650**	**$443,920**	**$491,770**	**5.2%**

Source: Gartner Dataquest (November 2001)

TABLE 3

Worldwide Telecommunications Services Market Revenue by Region, 1999–2005 (Millions of U.S. Dollars)

	1999	2000	2001	2002	2003	2004	2005	CAGR 2000-2005
Asia/Pacific	$191,008	$217,226	$246,369	$266,872	$288,371	$311,538	$338,805	9.3%
Central and Eastern Europe	$24,384	$28,445	$34,244	$39,349	$43,499	$47,878	$51,284	12.5%
Latin America	$78,535	$90,572	$105,708	$119,074	$133,481	$152,473	$167,954	13.1%
Middle East and Africa	$38,651	$43,057	$48,863	$55,817	$64,804	$75,879	$90,524	16.0%
North America	$267,410	$302,416	$338,365	$379,461	$422,470	$461,088	$494,591	10.3%
Western Europe	$204,515	$210,426	$235,588	$258,360	$278,361	$296,655	$313,178	8.3%
Worldwide	**$804,502**	**$892,142**	**$1,009,137**	**$1,118,933**	**$1,230,985**	**$1,345,512**	**$1,456,336**	**10.3%**

Source: Gartner Dataquest (November 2001)

help placate unhappy shareholders by postponing further investment.

- Telecommunications equipment vendors that had geared up to meet the huge increase in demand indicated by their clients were caught with excess costs when the rapid increase in demand failed to materialize and their stock value plummeted.

The economic downturn, which began to be recognized in the United States toward the end of the fourth quarter of 2000 and has spread to many regions since, has simply added to telecommunications infrastructure vendors' problems. This has further depressed general expectations for growth and stock-market valuations.

In the enterprise and mobile terminal sectors, the impact of the economic recession was perhaps more direct—both businesses and consumers delayed purchases and looked to reduce spending. Purchases of telecommunications equipment have undoubtedly been affected.

The terrorist attacks in the United States on September 11 may have a mixed impact on the telecommunications market. The events have definitely created further recessionary pressure, extending the likely period before economic recovery can begin but potentially leading to stronger growth toward the end of the period of economic recession.[7]

Some telecommunications service providers are expected to delay investments further while they consider options in the light of evolving customer requirements. On the other hand, telecommunications equipment suppliers could even benefit from increased user demand for more secure, more redundant, and more geographically diverse networks and services that the terrorist attacks may stimulate.

Reasons to Be Cheerful: The Bottom Line

It is important to recognize that the aforementioned negative factors will mainly affect the equipment markets. The telecommunications services sector, if not totally immune, is certainly less susceptible to this economic slowdown. The innate demand for basic voice capacity in developing markets, and higher bandwidth in all markets, remains. It is much easier for a company to delay a capital expense decisions (for both business users and consumers) than to reduce reliance on a service that has become part of its business process. Also, the bandwidth and connectivity implications of Web-based services and applications will drive growth significantly during the forecast period, even if this growth is slower than some of the dot-coms would have had us believe.

Even the excess in capacity in the backbone will not last long as usage continues to grow, while there are still shortages elsewhere in the networks. Many operators are taking advantage of reduced competitive pressure to hone their plans for the next big push. While there are clear opportunities today, a number of next-generation network technology segments may show relatively slow growth for the next 18 months but will pick up significantly after 2003.

The enterprise equipment and applications sector is also evolving, with converged networks being the major long-term opportunity. In the meantime, customers cannot delay system-upgrade decisions forever: wireless local-area networks (WLANs), IVR systems, and applications that help improve enterprise efficiency are growth opportunities.

The mobile handset sector may be experiencing its first real downturn, but infrastructure investment—increasingly in 3G—is still rising. From 2002, next-generation phones should start to have a positive impact on volumes in key markets.

Overall, the explosive growth across many sectors in 2000 was clearly unsustainable. Gartner Dataquest expects that the telecommunications market will not return to the growth seen in 2000, but it will return to growth.

Notes

1. For more information, see Gartner Dataquest's "Steady Demand Will Sustain Global Telecommunications Services Market" (TELC-WW-DP-0126).
2. For more information, see Gartner Dataquest's "Enterprises Refuse to Buckle Under Terrorism Threat: Focus on Business Fundamentals and Telecommunications Spend" (TELC-WW-DP-0105).
3. For more information, see Gartner Dataquest's "Mobile Markets Cool from White-Hot to Just Hot" (TELC-WW-DP-0123).
4. For more information, see Gartner Dataquest's "Service Providers Will Continue Cutting Capital Expenditure" (TELC-WW-DP-0128)
5. For more information, see Gartner Dataquest's "Mobile Markets Cool from White-Hot to Just Hot" (TELC-WW-DP-0123).
6. For more information, see Gartner Dataquest's "Enterprises Will Need More Bandwidth and New Systems" (TELC-WW-DP-0125).
7. For more information, see Gartner Dataquest's "Global Economy Faces 2001–2002 Recession" (HARD-WW-DP-0134).

The Softswitch and the Besiegement of Telecom Incumbency

Elon Ganor
Chairman and Chief Executive Officer
VocalTec Communications, Ltd.

Arnold Englander
Vice President, Product Strategy and Planning
VocalTec Communications, Ltd.

Introduction: Incumbency Besieged

During the past two years, much has been written about softswitches from technology, network architecture, and vendor/product perspectives. Numerous articles and conference presentations have defined and argued the benefits of softswitches as the key to converged data and voice networks, to reductions in carrier capital expenditures (CAPEX) and operational expenditures (OPEX), and to the rapid, flexible, and cost-effective creation and delivery of enhanced services.

This paper disputes little of that. In fact, as a vendor of softswitch-based services platforms to carriers worldwide, VocalTec has contributed many such articles and has made many presentations of that kind.

This paper addresses how and why softswitch-based services platforms are excellent competitive weapons for the telecom industry battles being fueled by worldwide regulatory changes.

Looking at the telecom industry today, we see a redefinition of the business of established carriers and the entry of many types of newcomers. We see local carriers moving into the long-distance business, and long-distance carriers moving into local services. We see national carriers setting up operations in other parts of the world, and we see companies that were not previously telecom providers (such as electric-power providers) now becoming telecom carriers as well.

Overall, we see a multitude of challenges to a historically fundamental phenomenon in the telecom industry—incumbency.

The markets of incumbent local-exchange carriers (ILECs) are being attacked (or at least their high-value markets are being "cream skimmed") by competitive local-exchange carriers (CLECs). Local-exchange carriers (LECs) in general are girding for their markets coming under attack by what have historically been interexchange carriers (IXCs). Relatively new global managed services (such as XA Alliance's pan-Asian telecommunication service network) are winning markets and traffic out from under from incumbent carriers. The primary (and often the former monopolistic) carrier of one country is now setting up its own facilities-based operations in other countries, attacking the market of those other countries' incumbents (often former monopolies themselves). A case in point is Deutsche Telekom in the Far East and South America. One may think of this phenomenon as incumbents, acting outside of their region of incumbency, attacking incumbents defending their region of incumbency.

And to make this even more complicated, the globalization of business is straining the network and business models upon which carrier-hosted services—such as Centrex, bridge conferencing, and voice virtual private networks (VPNs)—were originally built. None of these make sense any more as "local" or "long-distance" services (if they ever did). In today's global economy, these services need to be both local and long-distance—while also being relatively distance-insensitive in their pricing. If these services were "natural" for either IXCs or LECs (because they were implemented using feature sets on either Class-5 or Class-4 switches), this "birthright" or predestined incumbency is being challenged by real business requirements. The requirement today is for globally hosted business services. These create new opportunities as to who hosts them, who sells them, and the types of facilities off of which they are served. Hosted services need to be provided on a global basis regardless of the country in which a business has their headquarters, divisions, branch offices, or "road warriors"—one of the more complex challenges to the notion of telecom incumbency.

It is noteworthy that while the besieged are easy to characterize (incumbents and often former monopolies), the attackers are not. The attackers include well-established companies (including incumbents elsewhere that are now stepping out of their own fortresses), start-ups, converged network "next-gens," and companies large and small entering telecom from some other business, such as electric power.

This sets the stage—or battlefield—for explaining why we believe that softswitch-based service platforms are powerful siege weapons for attacking the fortresses of incumbents operating in their historical region of incumbency.

Enter Softswitch

Incumbency in the telecom industry has come under siege as a result of country-by-country telecom liberalization, deregulation, privatization, and competition. This has taken place rather suddenly (by telecom industry standards) after a 50-year (in some countries a 100-year) status quo. This has raised a Hydra of countless simultaneous issues for which most telecom executives were hardly prepared.

This process spread rapidly from the U.S. Telecommunications Act of 1996, through telecom reforms in each of 27 European countries during the second half of the 1990s, to India's New Telecom Policy of 1999 (NTP-99). Today, this process, and the industry re-alignment that it is causing, is far from over.

The term "softswitch" started to appear in the late 1990s referring to a type of switch that would replace central office (CO) circuit switches used in the traditional public switched telephone network (PSTN). The original motivation and advantages intended for softswitches included the following:

- Supporting voice-over-packet networks in addition to (or instead of) circuit-switched networks
- Supporting open standards to enable multivendor, best-of-breed networks and services platforms
- Supporting the use of industry standard hardware
- Supporting service-creation environments (SCEs) to dramatically shorten and reduce the cost of developing, trailing, and deploying new services
- Supporting multiple telecom protocols, each best for its own purposes and functions but requiring integration and inter-working
- Enabling lower cost and smaller footprint, but scalable, facilities

The Softswitch Consortium, an international organization, was formed in May 1999. A large number of vendor announcements—some premature—took place at Jeff Pulver's VON (Voice on the Net) conference in the fall of 2000, and more mature softswitches have been announced, trailed, and deployed since.

It is not entirely clear whether or not the idea of the softswitch sprang from telecom reform and the besiegement of telecom incumbency. It is clear to us, however, that as early as October 2000, several of VocalTec's customers (large and small, former monopolies and start-ups) were starting to appreciate that one or more aspects of softswitches could help them attack telecom incumbents in their regions of incumbency.

The connection between softswitches and the besiegement of incumbents is based on the following points:

- Softswitch-based service platforms can provide services that are location-independent and largely distance-insensitive. Services in Tokyo, Beijing, and New York can be served off of a platform located in Berlin, and calls between Tokyo and New York can cost almost as little to the carrier as calls between New York and New Haven.
- Global hosted business services served off of softswitch-based platforms are neither inherently local or long distance—they are both. The incumbent local carrier, with its Class-5 switch functionality, has no inherent advantage in providing Centrex services to the likes of Mercedes-Benz offices around the world; nor does the incumbent long-distance carrier, with its Class-4 functionality, have any advantage in providing voice virtual networks to IBM campuses in Martlesham, Boca Raton, and White Plains.
- Softswitch-based service platforms can be granular and scalable. This means that the cost of entering new markets—new geographically and/or in terms of types of services—can be very much lower than it has been in the circuit-based PSTN world. It also means that carriers can grow their networks, platforms, and service capabilities as customers and revenues materialize in these new markets. This goes a long way toward undermining the advantages of an incumbent having provided services to a region for the past 50 years, especially when telecom reforms encourage the besiegers to adopt selective or "cream skimming" service availability strategies from which the incumbents may be precluded.
- More and more, softswitch-based service platforms are supporting SCEs that enable rapid development, customization, branding, and fine tuning of value-added services by means of visual programming languages and Web development tools that are understood by a large community of developers (very different from the scarce resources needed to program new serves on a "hard switch"). This means that softswitch-equipped carriers can develop new services and new bundles of services that are custom tailored to target markets—and do so many times faster than incumbents can develop generic services on their "hard switches."
- Softswitch-based platforms can support converged voice and data services that cannot be supported well (in some cases at all) using the circuit-based "hard switches" deployed by incumbents during the last 10, or even 20, years. Some of these services are high-margin value-added personal computer (PC)–centric applications. Others simply involve the consolidation of historically separate voice and data operational costs, and connectivity, both on the carriers' side and on their customers' side, that lead to important cost savings. Either way, this convergence of data and voice services is available more readily and at lower cost to the softswitch-equipped besiegers than the "hard switch"–equipped incumbents they are attacking.

The Besiegers, Their Objectives, and Their Requirements

As a vendor of softswitch-based solutions and one of the major providers of IP telephony equipment worldwide, VocalTec Communications' platforms help its customers grow their businesses in ways that gain advantage from telecom reform, often by besieging incumbents in their regions of incumbency. The besiegers are start-ups, converged network "next gens," or incumbents (elsewhere) expanding beyond their region of incumbency. These carriers are growing their businesses in a world shaped by telecom reform. This is a world where the distinctions between IXC versus LEC, long distance versus local, and wholesale versus retail, are no longer clear. This a world that gives advantage to carriers deploying location independent platforms, providing distance-insensitive and globally hosted services, incurring fewer risks in entering new markets and selling to global business customers.

In particular, these customers tend to be the following:

- Packet telephony–based wholesale long-distance (LD) carriers and "carriers' carriers" that provide international interconnect to affiliate carriers
- Retail LD managed communications service providers serving large global businesses
- Large CLECs and LD carriers expanding into CLEC–like business models that focus on high-value business services and require the following:
  - The ability to provide bundles of business services (including voice VPNs [V–VPNs], conferencing, call centers, Centrex, voice-mail, unified messaging, and PC communications) on a global—thus LD as well as local—basis
  - Layered service-management systems for multi-tiered service customization, personalization, and self-provisioning
  - Low entry costs yet a high degree of scalability, both of which are important for carriers entering new markets, and especially to support "cream skimming" strategies

VocalTec softswitch-based solutions are specially designed to enable carriers in Category #1 to expand into Categories #2 and #3 and to enable carriers in Category #3 to differentiate their local business services from ILECs by satisfying their customers' global as well as local needs. VocalTec solutions are designed to enable carriers to compete with ILECs by building out service footprints according to a very focused, risk-contained "cream skimming" strategies. This requires integrated service platforms that are both fine grained and scalable (for low entry costs and fast growth as revenues grow) and that provide SCEs that support the rapid development of highly customizable services and service bundles that are simultaneously global and local in nature.

Such softswitches must meet the following market requirements:

- Rapid and easy multiservice integration, whether the features of the service are served off VocalTec or VocalTec partners' feature and media servers.
- Support for hosted business services (including Centrex and voice-mail) on a global (enterprise-wide or supply-chain–wide) basis, rather than on just a campus-wide or city-wide (CO or NXX local-exchange area) basis.
- Parity and uniformity of subscriber service features with media gateway control protocol (MGCP) or session initiation protocol (SIP) media gateways or end points.
- Multilayered service provisioning, customization, and personalization. This facilitates and reduces the costs of services resale, selling customizable services to businesses, and customer self-provisioning and personalization.
- Inter-carrier security and address translation. This enables carriers to interconnect with affiliate service providers using unmanaged networks, including the open Internet.

FIGURE 1

Softswitch and the Besiegement of Telecom Incumbency

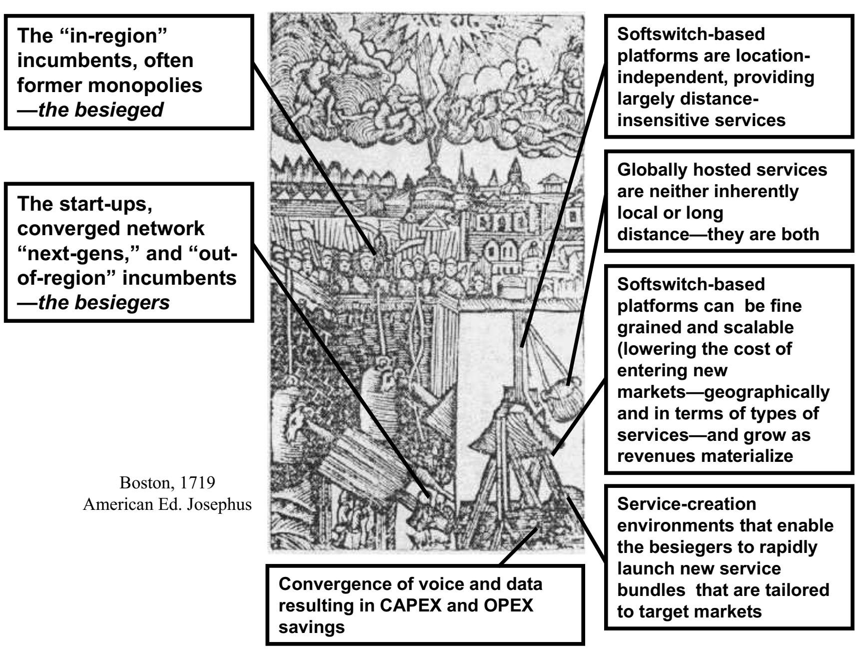

Boston, 1719
American Ed. Josephus

- Easy integration with already deployed H.323 voice over Internet protocol (VoIP) networks.

Conclusion: The Softswitch as a Siege Engine in Carrier Battles Fueled by Telecom Reform

We have portrayed the current state of the telecom industry as a "besiegement of incumbency" fueled by telecom reform worldwide.

We described how the softswitch is a powerful siege engine for start-ups, converged network "next gens," CLECs, and incumbents (elsewhere) expanding beyond their region of incumbency as they attack markets of "in-region" incumbents.

We have pointed out that a "region of incumbency" can be literally geographic and/or "horizontal," such as LD versus local, and voice versus data versus the convergence of voice and data.

We have also discussed how some softswitches—because of their granularity, scalability, and location-independent and relatively distance-insensitive natures—lend themselves to "cream skimming" strategies aimed at high-value user groups such as the global Fortune 1000 and their supply chains. We contrasted this with the "in-region" incumbents' having to juggle the needs and profitability of metropolitan, campus, suburban, and rural customers in their local-exchange area, in many cases using their "one-size-(must)-fit-all," inflexible, legacy "hard switches."

We also discussed the advantage provided by softswitches that enable trunk replacement, carrier interconnect, and carrier exchange services on the same integrated platform that provides service bundles that include V–VPNs, Centrex, call centers, conferencing, calling cards, voice-mail, unified messaging, and PC communications on a global as well as local basis.

The softswitch-based platforms are remarkably well-suited to a telecom industry landscape in which the boundaries between wholesale and retail, LD and local, trunking and subscriber, telecom infants and octogenarians, and voice and data are all embattled front lines in the besiegement of incumbency.

What remains to been seen is the speed with which start-ups, "next gens," and "out-of-region" incumbents—the besiegers—will use softswitches as siege engines, and the speed with which the "in-region" incumbents will find ways to use softswitches to fight back, both in their regions of incumbency and by expanding their businesses elsewhere—thereby becoming besiegers themselves (see *Figure 1*).

The Troubled CLEC: Causes, Symptoms, and Solutions

A.J. Germek
Chief Operating Officer and Co-Founder
GuideComm

Introduction

The Telecommunications Act of 1996 dramatically changed the topography of the telecommunications landscape. The legislation almost single-handedly helped to transform this once-staid industry from one dominated by a small number of large firms into a highly competitive market area. The 1999 projection of a $100-billion local-services market[1] set off a rush to finance the development of the infrastructure required to realize this potential. Venture capitalists and private investors alone pumped $17.68 billion[2] into this market segment over a two-year period, spawning some 2,000 new competitive local-exchange carriers (CLEC). Most of these service providers experienced a period of rapid growth and skyrocketing market valuations. These results worked to reinforce the edict to build out the network and acquire customers, both at the expense of deploying the robust back-office operational support systems (OSS) necessary to run the business efficiently.

Subsequently, the latter part of 2000 and the beginning of 2001 brought several high-profile CLEC failures, and the stock market as a whole turned downward, with technology stocks suffering in particular. The competitive service-provider market has since cooled considerably, and many service providers, reviewing the expansion of the past five years, are now reconsidering their business plans. Moody's Investor Service warned investors in March of this year to watch for credit-rating downgrades and defaults for CLECs and expressed concern that some of these companies will have trouble paying interest on several billion dollars of zero-coupon bonds when those bonds start to require cash interest payments over the next two years.[3]

The market opportunity for competitive communication services still exists; however, CLECs have captured just 6.7 percent of the overall market, a relatively small percentage given the supposedly ripe market opportunity.[4] But as Neil Sheh, vice president of M/C Venture Partners says, "CLECs are still great business. They just need to be built properly."[5] An increased emphasis on OSS is the key component to the CLEC turnaround strategy, offering hope to financiers who want to recover their investments, and giving the CLECs themselves renewed hope of capturing an enormous, as yet untapped market.

Market Conditions

Winstar Communications, GST Telecom, e.spire Communications, ICG Communications, IAXIS, NETtel, Star Telecommunications, and NorthPoint Communications are just some of the service providers who declared bankruptcy within the past several months. Others, such as Teligent, are reporting serious financial difficulties. How did the industry get to this point? There are various theories and opinions, many pointing to the overheating of the high-tech market and the vulnerability of the dot-coms. However, many of the problems can be traced to conditions and causes specific to the CLEC industry.

Root Causes and OSS Impacts

A variety of significant problems have combined to create the perilous state in which many CLECs now find themselves. For each of these root causes, an overlooked OSS has contributed to or magnified the problem. Conversely, a sharpened focus on the OSS will deliver improvements quickly.

Management Team Composition
As with any industry, the success or failure of any enterprise, especially a start-up, is directly related to the vision, skills, and expertise of the management team and their ability to execute. With the explosion of the CLEC market and demand greatly exceeding supply, the management teams of many CLECs have included an uncommonly high proportion of first-time executives. CLECs have also experienced management-team turnover, fueled by richer opportunities elsewhere or lackluster performance. As a result, developing and executing a consistent business plan represents a significant challenge to CLECs. Market and service plans are continually altered, and the OSS and information systems (IS) strategies that support them are also in flux, producing an unstable environment and systems whose potential is never fully realized.

Overexpansion
An investment environment where money was readily available, combined with the huge market opportunity, fostered an "If you build it, they will come" atmosphere. The prevailing wisdom among service providers is that long-term survival is

dependent on increased size and scale, and that being first to market creates a permanent advantage.

These factors created a single-minded focus to build out the network. However, this focus frequently caused CLECs to outrun their logistics and their OSS, which remained incomplete or in a state of operational infancy. Provisioning, quality of service (QoS), and even invoicing became challenging assignments for the IS and operations departments. Systems strategies were often ill defined or incomplete. The systems themselves required substantial manual intervention to complete the process of taking an order, provisioning/activating the order, and billing for services rendered. Substandard OSSs produced inefficiencies that greatly weakened the CLECs' competitive position. Or as Royce J. Holland, chairman and chief executive officer (CEO) of Allegience Telecom, stated, "I think it has been established over the last three years that in order to be successful, you have to have a state-of-the-art OSS"[6].

Competition from the ILECs

The incumbent local-exchange carriers (ILECs) answered the challenge of the burgeoning CLEC industry by using a combination of financial muscle and marketing strategy.

Competition for control of the market spawned price wars. The oversupply of basic services reduced prices to the point where CLECs could neither make a profit nor recoup their investments. For example, the average price of a T1 circuit dropped from $2,500 per month in 1998[7] to $1,296 per month in 2000[8]. Service providers are offering long distance for as little as $0.03[9] per minute, whereas the cost for that same service was $0.28 per minute just two years ago[10]. Overall, revenues per minute continue to decline as usage minutes continue to grow, putting a premium on efficient network utilization[11]. To differentiate themselves, service providers must offer bundled packages of additional services for that same price, thereby increasing overall delivery cost.

CLECs might have effectively countered the price competition presented by the ILECs by offering rapid provisioning or QoS guarantees, each far more feasible for the CLECs building new networks than for the ILECs encumbered with legacy systems. However, because the CLECs lacked the OSS performance necessary to deliver on these guarantees, they had no choice but to shift their marketing focus to price competition, playing directly to the strength of the ILECs. Meanwhile, the ILECs were able to target their marketing campaigns toward the familiar values associated with these well-known entities: size, experience, and permanence. These exploited the CLEC weaknesses in performance, speed, and agilityæ attributes that would have been strengths with a solid OSS foundation in place.

Voice Technology

Most of the CLEC buildout was originally implemented with voice technology or circuit-switching equipment, which in today's market is quickly becoming a commodity or complementary feature associated with data services. Hardware manufacturers that financed equipment purchases at 125 percent of cost or greater were delivering huge quantities of circuit-switching equipment into the CLEC market, along with additional capital for operations. CLEC and ILEC networks and their service offerings were virtually indistinguishable. By investing the additional capital more heavily into their OSS, CLECs might have offered services and features that would have differentiated them from their competitors, avoiding strict price competition. The net effect, however, is that carriers are faced with financing notes coming due, including approximately $150 billion in high-yield securities[12] (i.e., junk bonds); and in most cases, either their OSSs are not mature enough to provide a competitive advantage, or their network infrastructure is not suited to offer the data services that represent the future of the industry.

Dilution of IS and Operations Talent

The technical talent within the communications industry was diluted even more severely than the executive talent by the great influx of new players. Within the IS and the operations departments, many of the personnel placed in mission-critical positions came from the ILEC environment, where they only had experience with a portion of an OSS or had no experience with building systems from the ground up on a priority and rapid-deployment basis. As a result, large investments often produced ineffective results, leading to staff and budget reductions for ongoing operations.

The Penalties

Stock market prices have dropped below the level of a reasonable correction, and valuations are reaching new depths more alarming than the previous unprecedented heights. The average CLEC stock price is down more than 83 percent from its high in 2000.[13] For example, Covad Communications' stock price plummeted from a high of $67 per share in March 2000 to a value of $1.20 per share as of April 12, 2001. Teligent's stock price, $100 per share in March 2000, was down to $0.48 per share on April 12, 2001.

This financial performance has choked off the traditional funding sources for the capital-intensive service-provider business. With the stock market pummeling publicly traded firms, new initial public offerings (IPO) in the service-provider space have all but ceased. Venture capitalists, who view the IPO as their investment liquidation strategy, are slowing the pace of investment and raising the hurdles for new ventures. "If they [CLECs] are going to raise money, they are going to have to show some return on investment," said Yankee Group analyst Rob Rich.[14] Bob Nicholson, a partner at Spectrum Equity Investors, predicted in September 2000 that 5 to 10 publicly traded CLECs that were holding high-yield bonds could default by the end of 2000.[15] He was not too far off his prediction. And as the debt grows, debt holders may become the owners of many communication companies, a role they are neither suited for nor interested in assuming.

The Business Challenges Facing CLECs

In spite of everything happening in the market, the CLECs must face and overcome the following challenges:

- CLECs must reassert their ability to move more quickly and nimbly than the established ILECs in providing services.

- Revenue streams must become more reliable and robust and not simply be executed as proofs of concept.
- Cash-burn rates must be carefully managed or reduced as access to capital markets continues to diminish.
- Issues regarding access to capital notwithstanding, network expansion must continue at a measured pace to achieve size and scale.
- OSSs must be right-sized and effective to deliver the promised revenues and levels of service.
- Low morale must be checked through vision and decisive leadership at the executive-management level.

OSS: Keeping Pace?

Attention to OSSs (or the lack thereof) stands in stark contrast to the enormous network buildout undertaken within the industry. In many cases, CLECs have ignored or underfunded their OSS and have not made effective OSS deployments one of their top priorities. The lack of efficiency and/or automation in the back office regularly leads to lost revenue and service issues. Manual processing and segregated management systems have resulted in excessively long provisioning times and error-filled orders. Inefficient diagnoses and resolutions of network problems have deteriorated customer relationships, producing excessive customer churn and delinquent accounts. Those issues notwithstanding, it has been CLECs' inabilities to support new data products and services that have most dampened the enthusiasm of both customers and investorsæbecause the reality of the e-commerce/communications world is that data is paramount.

Offering new products and services efficiently and seamlessly will enable CLECs to grow and prosper. And a robust, flexible OSS is a critical component to realizing that strategy. XO Communications, a company based in Reston, Virginia, is one CLEC that has embraced this view. And XO is one of the few CLECs that has invested wisely and built a flexible, scalable OSS. In fact, current XO CEO Daniel Akerson, former president at MCI, credits "the best back-office systems in the industry" with helping MCI overtake AT&T in the long-distance marketplace[16]. And as the chief executive at XO, he is championing the development and implementation of a robust OSS that is contributing to this CLEC's growth through expanded products and services.

For instance, XO offers a service called XOptions, which allows a customer to decide their local, long-distance, data, and Web-hosting needs and to choose a plan that charges one flat rate. XO delivers all these services for 10 to 40 percent less than customers would find elsewhere, and because of their strong back-office system, it is all delivered on a single bill[17].

A well-defined OSS system strategy, workflows, architecture (open, flexible, and scalable), component prioritization, and software selection are critical to success. For example, if successfully implemented, an OSS with an automated order-entry and provisioning system can often provide the following types of results:

- One CLEC reported a 23-percent increase in the number of lines installed from the previous quarter (when the company was using manual processes) to the first quarter of using automated order entry and provisioning.[18]
- Another CLEC showed 62 percent of their orders being filled in less than 5 days, provisioning volume of up to 9,500 phone lines per day, and order fallout rates that dropped from 43 percent to less than 2 percent.[19]

In the rush to build a state-of-the-art and expansive "physical body" (the network), most service providers were not able to give that body the "brains" (the OSS) to operate efficiently and effectively. Only now are the penalties for this failure starting to emerge.

Measuring the Health of the OSS

In addition to the obvious problems of unacceptably long order and provisioning lead times, excessive customer complaints, and attrition, a CLEC must also be able to detect the more subtle signs of underlying distress within an OSS. To identify an ailing OSS, a CLEC should check for the following problems:

- Long IS development and implementation timeframes
- Missing or inadequate OSS components
- Rapid growth in the IS budget without obvious benefits
- Shifting timeframes in the delivery of OSS and IS projects
- Efforts are focused on stabilizing current OSS components rather than adding functionality or coverage
- Revenue growth constricted by inadequate/lack of OSS support
- Provisioning lead times approaching or exceeding the incumbent service providers
- New services cannot be delivered through the current OSS structure
- Unfulfilled requirements and unmet customer expectations for software-produced services
- Confusion or conflict as to OSS goals, objectives, and timeframes

The Solution: Short-Term, Tactical Initiatives within a Strategic Architecture

After spending hundreds of millions of dollars on the assets, people, and systems to manage the networks and services, a company must then decide where to start to address the problem.

The overall goal in improving the health of the OSS must be to meet the short-term, tactical business objectives without harming the long-term growth potential and architectural integrity. This requires a critical analysis of the current OSS position relative to the overall OSS landscape (as depicted in *Figure 1*) and a decisive, positive approach for the future. The analysis and OSS strategic planning must account for revenue growth, reduction of expenses, enhanced performance, specific business objectives, investment options, and identifiable metrics. A well-defined OSS strategy and architecture offers support across all levels of the CLEC operation.

The guiding principles should include the following:

- A focus on effective results in short, fixed timeframes
- The ability to scale when the business climate improves and growth becomes imperative

FIGURE 1

Well-Defined OSS Strategy and Architecture

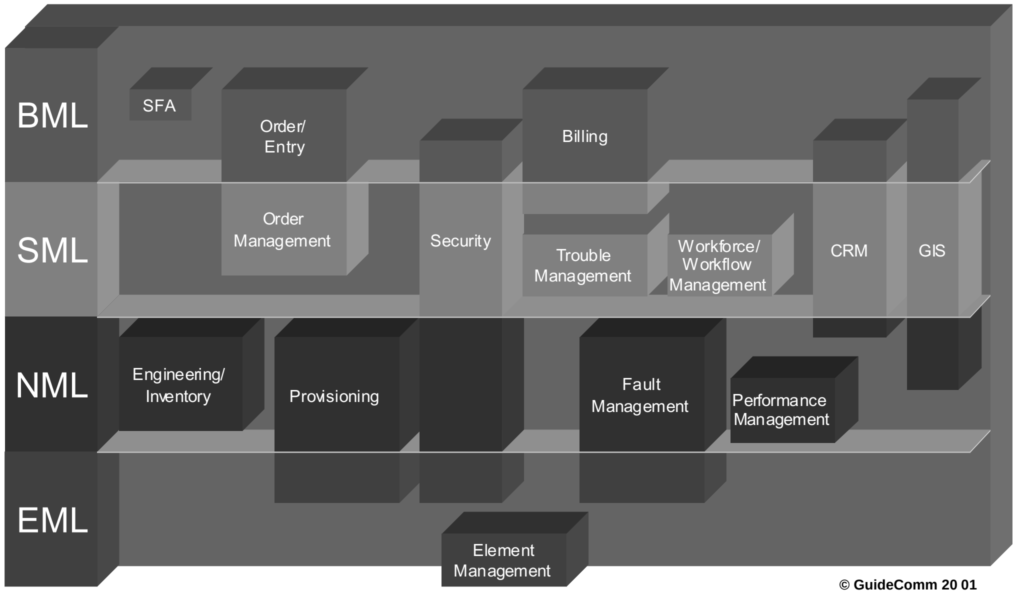

A targeted OSS analysis should produce an action plan that when executed provides a path to growth and profitability. An experienced OSS strategist can complete a targeted analysis within a 30- to 60-day period with an efficient, structured approach. The tasks to be executed as part of the analysis are illustrated in *Figure 2*.

After identifying the OSS implications and related business objectives, recommendations and alternatives for the OSS can include the following actions:

- Reprioritizing and rescaling markets and service priorities
- Implementing rapidly-installed new software components
- Reducing functional requirements to address critical and short-term needs
- Re-architecting OSS solutions for flexibility and scalability
- Eliminating non-essential, long-term functionality and projects
- Adding small-scale critical components and interfaces
- Determining functional outsourcing potential

Based on the analysis and recommendations, the CLEC can create and execute an action plan that includes the tasks and budget against which success can be tracked and measured.

Independent Review

Any department or executive may have a commitment to the current OSS status. Marketing may have committed to service offerings or products. Engineering may have committed to network buildouts and equipment purchases. Operations may have committed to processes and staffing levels. IS may have committed to a comprehensive OSS plan and strategy. These vested interests can inhibit a critical assessment of the current OSS status and the development of a strategic solution.

FIGURE 2

Targeted OSS Analysis Tasks

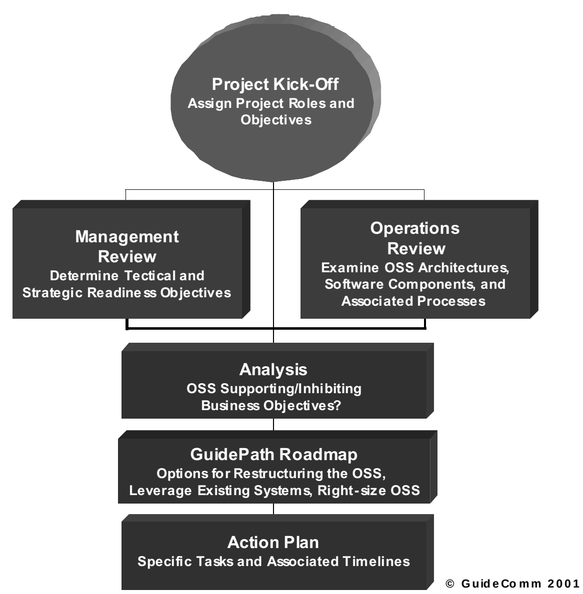

Therefore, the OSS analysis needs to be performed by a third-party OSS professional who is not vested in the current OSS strategy or predicament but is driven to achieve results. The expert guidance needed to solve this dilemma will require business acumen as well as OSS expertise and an independent review.

Finally, a willingness on the part of the CLEC to dedicate the resources and leadership necessary to execute the finalized action plan is essential for the health of the OSS. However, once the OSS is on track, the CLEC will be positioned for survival and leadership in the ever-growing, ever-competitive communications market.

Notes

1. Morgan Stanley Dean Witter: Global Investment Research, March 16, 1999
2. *The State of the Competition in the U.S. Local Telecommunications Marketplace: An Annual Report of ALTS,* February 2000
3. American Bankruptcy Institute, March 28, 2001
4. Economy.com, January 31, 2001
5. *TelOSSource Magazine*, July/August 2000
6. *TelOSSource Magazine*, July/August 2000
7. *Computer User*, December 20, 1998
8. Association of Telecommunications Managers, November 2000
9. Detroit News.com, January 1999
10. Morgan Stanley Dean Witter, April 2000
11. Association of Telecommunications Managers, November 2000
12. *Inter@ctive Week*, September 4, 2000
13. *Barrons*, November 27, 2000
14. *Inter@ctive Week*, May 22, 2000
15. *ZDNet Interactive Week*, September 4, 2000
16. *Barrons*, November 27, 2000
17. *Barrons*, November 27, 2000
18. Anderson Consulting Web site, November 2000
19. Anderson Consulting Web site, November 2000

Evolution of Technology versus Carrier Adoption

Howard J. Gunn
President
Entrusted Communication Services

Abstract

The evolution of computer and communications technology and the application of new technology to solve current business problems are similar to a biological process. For a variety of reasons, it is relatively easy for U.S. business visionaries to come together with technology innovators to create a new offspring. Likewise, there seems to be a ready supply of enterprise "early adopters" who can see the business value of the new applications. The early adopters in a market typically help the embryonic application fulfill its stated purpose, or they kill it off in favor of even better applications that solve the business problem.

Many innovators and early adopters perceive a major gap between their market-driven evolution of new technology and the carrier adoption process of the carrier industry. The author believes the adoption gap is fostered by the "regulatory economics" of the carrier market, versus the "profit" motives that drive the enterprise market. This paper explores the perceived "gap" and how the Telecommunication Act of 1996 has brought the gap into clear focus in the carrier market.

The current gap is big enough to generate new efforts to resolve the problem. Newly proposed solutions to the perceived gap range from structural separation of wholesale and retail services to a return to monopoly yesteryear with some members of Congress ready to diddle with what has been called "the worst set of compromises and pork barrels ever passed by Congress." Based on the regulatory track record of Congress and the States, the author believes that any additional legislation will tend only to widen the gap between technology evolution and carrier adoption. The author also suggests that another regulatory fumble could lead to even more legal and market confusion than the set of regulatory changes that have been in the courts since the 1996 passage of the Telecommunications Act.

The collapse of the dot-com market, followed by the collapse of the new data local-exchange carrier (DLEC) and competitive local-exchange carrier (CLEC) markets were not caused by the innovators and early adopters killing bad applications in the market. These collapses—followed by the collapse of the long-distance carriers, equipment suppliers, and growing pressure on the incumbent local-exchange carriers (ILECs) stock prices—are probably a part of the technology evolution gap, stimulated by fundamental changes in public policy and the inability of regulators to create a viable open market in the local monopoly.

The paper concludes that the carrier technology-adoption process is changing slowly. The traditional adoption rate was similar to the pace that mainframe-computing operators accepted for personal computing and servers. Over time, dinosaurs supporting centralized computing and common business-oriented language (COBOL) programming either died off or evolved new technologies to support the distributed control paradigm of personal computing. New service control points (SCPs) are the first positive signs of a change in carrier adoption strategies. Even after the death of 40 percent of the carrier market—lead by Nynex, Pacific Telesis, U S West, GTE, and numerous CLECs—few carriers have started to plan a new technology adoption strategy for upgrading or converging their costly rate-base–driven overlay infrastructure.

Background

Congress passed and the President signed the Telecommunications Act of 1996. Although some have criticized the law as "the worst set of compromises and pork barrels ever passed by Congress," the law clearly reflects a significant change in public policy toward the local carrier monopolies. Prior to 1996, the laws protected public carriers from the impact of open markets. The new law enabled regulatory changes to free consumers from the power of the local monopolies.

The law created a way to grant new local carrier franchises that enabled the creation of CLECs and DLECs. The law also generated a mountain of legal issues that extends all the way to the U.S. Supreme Court. Some challengers are trying to protect the status quo monopoly, while others seek protection from the monopoly introduction of various scorched-earth pricing and proprietary service models.

Clearly, the most controversial issues are associated with the legal elimination of the local-exchange monopoly. This change in regulatory policy includes the involuntary unbundling of the current monopoly infrastructure and the Federal Communications Commission (FCC) mandate to price the unbundled elements at wholesale to new carriers using forward-looking total service long-run incremental cost (TSLRIC) methods. The unbundling of service elements and the TSLRIC cost model strikes at the heart of the local monopoly—the rate base.

The rate base in laymen terms is the accumulation of investments used to value a company, as determined by the regulating agencies and not the market. Over many years, this arbitrary regulatory valuation process evolved as a group of companies that acted as quasi government agencies known as public utilities. These agencies operated on a regulated cost-plus basis in power, water, cable, and telephone services markets, as well as other smaller market segments.

What makes the U.S. utility experience unique is that public stockholders own these cost-plus, quasi-public agencies. Without comment as to legitimacy or the historic value processes, in all cases, the rate base was created by administrative government actions that included federal, state, and, in some cases, municipal government oversight. The rate base and resultant carrier pricing strategies were initially created to reduce the capital risks in capital-intensive markets. The various rate-base companies were also driven by social mandates of government. In the carrier space, the mandates included "Universal Service," a government-controlled fair rate of return on the rate base and an indirect tax collection process that hid certain taxes in the service prices.

In the late 1980s and early 1990s, a variety of government efforts were initiated, trying to unbundle the rate-base cost-plus processes. For a number of reasons, fair-rate-of-return rate-base regulation and its arbitrary cost models was shifted to a mixture of government direct collections (taxes), plus a return on investment (ROI) income cap. The ROI regulation process was introduced as a way to stimulate a change in "utility economics." It was thought that these changes in process would stimulate an upgrade of the infrastructure. These early efforts culminated in the Telecommunications Act passed by Congress and signed into law by President Clinton on February 8, 1996.

The involuntary unbundling of the bottleneck monopoly facilities and the granting of new local service franchises coupled with TSLRIC methods were initiatives to circumvent the rate-base issues on an economic plane. The TSLRIC approach implies that the value of the bottleneck facilities should be valued at the cost to reproduce them with current technologies and not the arbitrary "regulated valuation" created over the last hundred years.

As an example, assume a hammer as a carrier unit of unbundled capacity. Today, hammers sell for $20. The FCC mandated that the valuation of all similar hammers should be $20. (The cost of a unit of capacity is its estimated reproduction cost using today's technology.) This new valuation would be used to set the wholesale price for renting old hammers to new carpenters, until they could afford to buy their own hammers. The problem is that the rate base contains millions of hammers that on paper carry a residual regulated valuation. This arbitrary rate-base valuation is similar to the million-dollar hammers associated with many cost-plus contracts, where unit valuation and cost allocations are jockeyed to fit budgets and other external variables.

The problem of not using the current market price of a hammer for setting unbundled prices is just as complex as using the arbitrary regulated valuation of hammers. The traditional carrier suppliers are accustomed to developing products and pricing strategies that ship million-dollar hammers to cost-plus operators. Because of this distorted regulatory valuation process, the carrier suppliers have typically had limited product success in open markets.

The last new technology adoption processes in the regulated market were products introduced in the 1980s (hybrid analog-line and time-divided switches, synchronous optical network [SONET] and asynchronous transfer mode [ATM] transmission systems). These adoptions occurred long before distributed processing and the deregulation of the local exchange. Although the cost of the vintage technologies has declined from about approximately $600 per subscriber to $200 per subscriber, new hammers using current technologies are selling for less than $100 per hammer.

The reasons that carriers do not adopt the new lower-cost technologies at the rate of technical evolution are varied and complex. The rest of this paper explores some of the structural reasons for the technology adoption gap versus the new technology evolution rate.

The Stage Is Set

The relatively new technology companies that have created the $100 hammer technologies perceive a major gap between the introduction of their technology and the adoption of their technology by carriers. Some type of gap exists in their mind, as evidenced by the 12-year interval since new technology adoptions by carriers. Many think the gap is caused by the difference in technology. Although this difference in technology could be a contributor, it is also reasonable to assume that the actual rate of adoption is a combination of the slow transition away from the cost-plus rate base, the confusion in the regulatory processes caused by the Act, some poor decision-making by management, and the current product process in the hammer business.

From the technology perspective, the stage is set for tremendous business growth. Today, 50 million people use data services. Today, only 22% of the U.S. enterprise market has high-speed access, and less than 2% of the mass market has high-speed access to data services. Given this stunning market situation, why has the carrier market collapsed? Why have major suppliers (Lucent, Nortel, Alcatel, Motorola, etc.) and many service suppliers (WorldCom, A&TT, Sprint, Enron, Williams, Covad, Exodus, etc.) flirted with the penny stock range? Why are digital subscriber line (DSL) prices increasing when supply is expanding and the volume of purchases is increasing?

Would faster carrier adoption rates solve the supplier and carrier market problem? Would involuntary unbundling or the monopoly-proposed rebundling of data service be better ways to solve the problem? These are good questions that will be addressed during the next several years.

The best way to understand today's carrier market turmoil and potential solution is to review the regulated market. First, it is important to consider that incumbent monopoly carriers are not stupid. They are well versed in the art of controlling the availability of goods and production resources. This art of market control is the very nature of a monopoly. The incumbent operators clearly understand that one more day of monopoly is better for their stockholders than any deal the government will cut with them with respect to the involuntary unbundling of the monopoly bottleneck facilities. No one should be surprised that they are trying to rebundle data services, enter the long-distance service markets, and not unbundled current bottleneck elements.

From an incumbent perspective, delay is clearly a good management tactic for thwarting the will of Congress—with the hope that Congress or the courts will change public policy with respect to the dissolution of their monopoly power. Again, without any intent to pass judgment, it is clear that open-market concepts only make sense to those that do not have a monopoly.

In traditional open-market terms, monopolies are harmful. Their centralized control of production resources, capacity availability, and prices ultimately works to the detriment of consumers. The open market of the computing industry is often used as a paradigm comparison to monopoly.

The open-market computing industry has fostered an entrepreneurial environment that enables visionaries to come together with technologists to foster new technologies for solving real business problems. In fact, this open-market marriage of business and technology has become an almost biological process, fostering new offspring and creating new markets.

For a variety of reasons, in the U.S. open market, there exist a relatively large number of innovators and early adopters in the enterprise business segment. These early-stage buyers can see how the new hammers help them to solve their business problem and gain a competitive advantage. They rapidly adopt the advances and then push the carriers for ubiquity to reach outlying locations. It has been at this stage of creating ubiquity that the engine of new technology, running at high speed, runs into the carrier that does not seem to be moving at all.

The first signs of the pending collision between the open-market technology evolution and the monopoly carrier adoption strategies occurred in the late 1980s and early 1990s. Carriers were adopting centralized control mixed analog digital switches, SONET, and ATM. Consumers, on the other hand, were adopting personal computers (PCs), distributed processing, and local-area network (LAN) technologies. The debacle of integrated services digital networks (ISDNs) illuminated this divergence from the carrier-based centralized control systems.

The regulatory rate-base and monopoly logic seemed to be that if an analog telephone line is worth $30 and a private data line is worth $90, then two digital lines with voice and data service features and a signaling channel (2B+D) must be worth at least $120. In retrospect, the monopoly pricing of ISDN in the United States was seemingly established to control the new service adoption rate and minimize the analog-line replacements, at the expense of the new technology.

Even worse regulatory logic seemed to be applied to T1 services. If a point-to-point T1 line was worth $5,000, then a switch enabled Primary Rate line running on the same wire must be worth a lot more. Needless to say, the old hammer makers priced their new technology to the monopoly market reality. Traditional carriers effectively avoided the cost of the digital buildout by not only overpricing the technology, but also making it almost impossible to get. A T1 would take anywhere from 90 to 180 days to install, and the buyer in many cases had to pay for the buildout.

But, sound monopoly practice in a rate-base environment was being exposed to the light of a deregulating environment. In 1988, the FCC deregulated the local carrier private-line market. Newly authorized specialty carriers captured a large percentage of the Primary Rate and T1 market. Monopoly carriers responded with T1 price declines, trying to use monopoly price control to manage the new deregulated market. It didn't work—the prices and margins collapsed for the incumbents, and a significant number of buyers went with the new carriers anyhow.

The new hammer makers, though, where flabbergasted by the ISDN technology adoption results. First, it should be understood that most new hammer makers were coming from an entrepreneurial open-market background. Their experience implied that if they could provide twice as much communications capacity and remove 100% of the transmission impairments of analog lines on the same two-wire infrastructure, then everybody should buy a new hammer.

Although the visionaries and technologists achieved their technical goals—as evidenced by the delivery of the fastest growing data service of all time (Frame Relay Q.921 is part of the original ISDN specifications)—the adoption rate of ISDN carrier subscribers was abysmal. At best, about a million new Basic Rate lines were sold during a 10-year period at the prices quoted by the carriers.

The second open-market technology adoption collision occurred with the advent of the Internet in the mid-1990s. Again, visionaries and technologists, in cooperation with government, created an exciting new medium for the rapid exchange of data. Some 50 million subscribers bought computers and modems at the enterprise and consumer level to use this new medium of knowledge exchange. Nine thousand new Internet service providers (ISPs) sprung up to manage the modem pools that connected the analog telephone network to the new data infrastructure. Some 10 percent of the subscriber base bought a second analog line in the period from 1990 to 1998 for unfettered access to the new data services.

Carrier adoption of the distributed control Internet technology was almost nonexistent. Carriers sold the analog hammers to a screaming throng of data buyers. Carriers even sold some of the ISDN lines as private-line access to the ISPs. Few of the carriers became successful ISPs, where managing the new distributed server and modem pool technology was the key to success.

Although the carriers stayed out of the new distributed processing data market, the switching and transmission manufactures that made the hammers for the carriers saw a different picture. The new hammer makers where producing technology that could tentatively replace the current carrier technology. This confusion had significant impact on the ability to sell the old $600 hammers. The carrier suppliers responded by reducing prices and buying up many of the new hammer makers, under the expectation that their carrier customers would enter the data market. The suppliers hoped for the development of a replacement market for the embedded transport and switches. The carrier suppliers, like many on Wall Street, failed to consider the implication of replacement technology and regulatory rate-base economics.

The Irresistible Force Meets the Immovable Object

Carrier adoption strategies and the rate-base economics for replacement technologies do not match up well with the direction or the speed of new technology evolution or uses. In most cases, carrier adoption of new technology is driven by the monopoly/oligopoly desire for new revenue from an installed base of current subscribers. In these cases, carrier adoption means building a new technology or application-specific overlay (adding to the cost base) and marginal pricing of the new services to capture new revenues, in addition to the current service revenues.

What isn't obvious in this monopolistic process is that this strategy sets prices for the new technologies in such a way that the price limits demand to those who are willing to pay the incremental premium price on the cost-plus infrastructure assumed in the overlay.

It may not be obvious that the adding of a new technology as an overlay is not generally intended to optimize the carrier cost structure. It is also not obvious that carriers intentionally use the unoptimized infrastructure cost, plus the new technology overlay cost, to protect the current services revenues and ultimately delay the displacement of current or even obsolete technology (replacement). As an example, a typical carrier operates 15 to 20 different application-specific and, in most cases, technology-specific overlays. The impact of these overlays on service prices and the business strategy of the monopoly have become the immovable objects that drove the adoption gap and stimulated the worldwide effort to deregulate the utility markets.

Technology, driven by the enterprise early adaptors, is not typically focused on the generation of new revenue. Enterprise advances are more typically focused toward the re-engineering and ultimately the replacement of existing processes that constrain the enterprise. The enterprise adoption rate is usually driven by a business need to improve decision-making or reduce per-unit costs. Acceptance of technology that can increase speed of information exchange, facilitate data availability, train the organization, collect necessary data and/or minimize cost per unit of revenue are readily deployed to gain a competitive advantage.

The irresistible enterprise force of "better, cheaper, and faster" are currently evolving technologies that favor client control, distributed computing, peer interactions, personal customization, and remote servers—the antithesis of the traditional centralized control and service overlay that drives new overlay service revenues and the carrier adoption market. Moreover, the enterprise speed of adoption of distributed computing and new software is creating a relatively short duty cycle of three to five years. This cycle is in stark contrast to the carrier adoption and duty cycle of 20 to 30 years or more between waves of technology.

The Product Life Cycle and the Carrier Duty Cycle

The product life-cycle concept can be thought of as the process of a product moving from sellers to buyers in a type of orderly progression. Buyer types range from innovators to early adaptors, to early majority, to late majority, and to stragglers. Each phase in the product life cycle calls for unique strategies to deal with the different types of buyers in the various market segments.

The carrier duty cycle is not at all like the product life cycle. The duty cycle is the interval of time that a certain technology exists as a dominant force in a monopoly market. The forces of the open market do not impact the duty cycle. From a regulatory perspective, utilities were granted "monopoly" status to minimize the cost of capacity. By restricting the investments, a product could be designed for an extremely long useful life—independent of its economic life, depreciation life, tax life, market life, or financial life. Rotary phones that lasted for years were an example of this duty cycle concept. In carrier terms, the monopoly position enabled technologies to be deployed with 40-year duty cycles (power plants, telephone switches, sewer systems, etc.).

For a number of reasons, the power of personal computing has produced a three- to five-year duty cycle as the users progressed from Commodore, Apple, and x86 PCs to today's 1.2 gigahertz (GHz) or more supercomputer that sell for less than $2,000. Software applications have tracked this same short duty cycle, as evidenced by Windows 95, 98, ME, 2000, and XP.

This short duty cycle and rapid adoption rate has been tough to support in the computing market. The forces of technology adoption driven by this cycle has already bankrupted an entire generation of host computer suppliers, an entire generation of minicomputer suppliers, and even a high percentage of the PC suppliers, trying to keep up with the ever-expanding demand for personal computing resources and communications between these 50 million users.

In terms of carrier duty cycles, there have only been three waves of carrier transmission technology since 1886, when the phone was invented. The first wave was analog transmission. This technology maintained a 100% market-share until the mid-1960s, when asynchronous digital transmission systems were introduced. The third wave, synchronized digital transport called SONET/ATM has been shipping since the early 1990s.

Analog transport still dominates the carrier transmission market. Asynchronous digital is still being shipped (T1 is the primary service link to digital service, and T1 is not a SONET–based signal). SONET/ATM is still in its build-out phase, with limited use as a replacement, after some 12 years of production. For all practical purposes, the transport duty cycle and replacement market is limited to replacing failing copper cables and the displacement of backbone routes through the addition of new optical SONET/ATM lines that upgrade the embedded analog and asynchronous digital infrastructure.

Switching system duty cycles have also been extremely long. Switching has gone through five waves of technology since 1886. The first wave of switching was manual operators and cord boards. This wave started in the 1880s and lasted till 1978 or so, when the last cord-board exchange was retired. The second wave, mechanical switching, started in 1899, when Almon Stowger invented the first automatic switching machine. The mechanical switching wave lasted until the mid-1980s, when the last new start shipped from Automatic Electric. The analog stored program switching (SPC), or third wave, began in the early 1960s with the introduction of solid-state electronic switching fabrics with centralized computer-controlled calls. The analog switch wave isn't yet complete, some 40 years later. The fourth wave, or digital switch era, started in early 1976 with the 4E in Chicago and the TRW ITS 4 in Ora Loma.

This centralized digital switch wave has barely crested, 25 years later. New start digital tandem business was so brisk that both Cisco Systems and Lucent Technologies bought new technology digital tandem hammer makers within the last several years (Summa 4 and Excel).

The fifth wave, generally labeled fast packet switching, began in the early 1990s, with the introduction of ATM self-routing switch fabrics with optical line interfaces that are SONET compatible. This wave of centralized switches arrived in the middle of the transition from centralized hosted computers to distributed processing PCs. The switches also arrived in the middle of the transition from rate base to open systems.

From a technology standpoint, the irresistible force of distributed computing embraced by the distributed enterprise ran into the immoveable object, the rate base and duty cycle, consisting of host controlled central switching and dumb transport systems that dominate the carrier networks. While the enterprise was rejecting ATM for its purposes, carriers were installing it for theirs.

It became apparent in the late 1990s that a technology gap existed. Why it existed, why it was widening, and how to close it seemed to confuse new enterprises and Wall Street. In most cases, this confusion is reflected in the financial results of the start-ups entering the carrier market.

As, an example, new enterprises operating as CLECs embraced the old monopoly carrier technologies, built multiple overlays without a rate-base–driven revenue stream, and went bankrupt. New enterprises operating as ISPs built distributed computing centers and went bankrupt. New enterprises operating as interexchange carriers (IXCs) and MNOs dug holes, filled them with sand, and went bankrupt. None of the bankruptcies were caused by a failure of the technologies or the lack of customers for the new enterprises. They failed in their fight with the incumbent monopoly.

The only new enterprises that seem to have survived built their new carrier business on resold lines and SCP software. Prepaid long distance, prepaid local service, prepaid wireless, and a variety of caller ID, call centers, and 800 services have come to market and thrived during the last five years.

Illusion or Conclusion

As discussed in this paper, the evolution of technology and the carrier adoption cycles have very little in common. Technology evolution is driven by a buyer's need to improve a business process and gain a competitive advantage through early adoption. Carrier adoption is driven by a need to increase revenues from existing customers. Because the adoptions are targeted at premium market segments, adoption is typically based on simple overlays that are priced to avoid or minimize change in the current revenue flows. Based on the current realities of the carrier adoption process, any new technology evolution must lend itself to the overlay market, not the convergence market. The successful evolution and adoption of SCPs by a number of carriers is an example of the process.

As described in this paper, incumbent carriers with a large installed base of paying subscribers and monopoly control over their market gain virtually no benefit from adopting new or replacing old technologies. As such, synergy and convergence terms in the enterprise market have virtually no benefit for the carriers. In fact, because of equipment duty cycles, carriers benefit the most from the adoption of incompatible application-specific overlays and selling new lines at higher prices.

It should be clear that the author is trying only to explain the ambiguity between technology evolution and carrier adoption. The paper is not intended to be an advocacy for or against the involuntary unbundling of the local exchange or the adoption of processes to circumvent the rate-base economics. The paper is simply trying to point out that monopolies do not benefit from unbundling or the creation of a wholesale local market. Given the relationship with stockholders, it makes sense for the carriers to present roadblocks, create obstacles, and develop pricing plans that ultimately disadvantage the CLECs.

Lastly, the paper tries to describe the difference between product life cycles and the carrier's view of technology duty cycles. Given that duty cycles are fostered by the rate-base economics, it seems highly probable that most carriers have not even begun to consider any next-wave adoption plan or replacement technology.

Integrated Access Devices: Bridging the Gap between Modern IP Networks and End Users

Peter Hall
Chief Operating Officer
Braintree Communications Pty Ltd

Integrated access devices (IADs) and intelligent protocol conversion is nothing new. However, the importance of the two is increasing rapidly as new-generation Internet protocol (IP) networks become more readily available for small-to-medium enterprises (SMEs).

The importance of IADs is rapidly increasing due to the slow uptake of modern IP networks by the SME market. Modern networks offer numerous added benefits to SMEs, so it is difficult to understand why the migration to the modern networks is slow. To explore the issues facing the SME market that is influencing the slow uptake, we will discuss the role that IADs and protocol conversion can play in overcoming this problem.

There are many facets to the SME markets reluctance whilst home users and larger corporates are experiencing the benefits that modern networks have to offer. Why is it that the SMEs maintain their existing legacy networks?

To understand the situation, we need to look at all stakeholders involved with migrating SME end users onto IP networks. The stakeholders exist of three major parties in the communications process: the carriers, the service providers, and the end users.

To understand the end-user environment and reluctance to migrate, business needs to understand the end user's perspective in migrating their businesses to a new network infrastructure. Do the benefits of the new network outweigh the costs involved in accessing this network?

The SME market does not argue that modern networks offer more benefits in comparison to their existing infrastructure. SMEs can gain access to new services, faster connection times, more reliable services, and long-term cost reductions. If SMEs can easily relate to the benefits, then why not migrate to the new networks?

There are three hurdles associated with migrating to modern networks facing SMEs, with one common underlying theme. The three hurdles consist of high capital costs in migrating, higher employee training costs, and higher support costs.

The underlying theme is the increase in costs in migrating to modern networks. Although SMEs can understand the benefits associated with the new networks, the costs associated with migrating are far too high for the majority of SMEs. SMEs do not have access to large amounts of disposable capital that would enable them to migrate. SMEs cannot justify migrating to the elite option if their existing basic network infrastructure performs the job adequately and reliably.

To overcome such reluctance, carriers and service providers need to address the end-user hurdles. The first financial hurdle is to replace all existing legacy equipment with modern appliances. Modern network protocols are significantly different from many legacy network protocols. Customer-premises equipment (CPE) needs to speak the correct language (protocol) to be able to migrate to the new network. Therefore, SMEs will need to replace all existing equipment. This is simply an unrealistic financial burden that many SMEs cannot justify. There lies the largest hurdle in the reluctance to uptake new networks.

The second hurdle is the increased costs in employee training. All employees will need training on the new equipment, which will not only be expensive, but also slow and frustrating for the SME. Most SMEs may need to replace all existing equipment simultaneously, not allowing the migration to be phased in over time. This places a significant financial burden on SMEs through employee training and business downtime.

The third hurdle is the support costs. Learning on the new equipment SMEs will no doubt require ongoing support as they adapt their new network to "fit" their business processes.

After understanding the situation from the end user's point of view, we can now look at the problems facing the other stakeholders in the communications process. Business needs to understand the importance of migrating existing customers to modern IP networks.

As modern IP networks evolve, it is becoming increasingly important for the carriers and service providers of communications networks to migrate existing customers from redundant legacy networks onto these modern networks. The modern IP networks consist of numerous cost-saving benefits for the carriers and service providers, particularly in network management and maintenance.

Carriers and service providers also have limited services that they can offer to existing customers on these legacy networks. New services that are being developed are increasingly designed for the modern IP networks with the aim to phase out costly legacy networks.

However, carriers and service providers do not want to tell customers throw to out all of their old hardware because there is a new network. They are faced with competitive discounting to entice customers onto the new networks.

Rather than trying to sell new products that are compatible with modern IP networks, carriers and service providers need to consider the SME markets that do not have access to disposable capital or the business case to upgrade their existing equipment.

When looking at what is connected to existing networks, there is a big list of legacy protocols. Carriers and service providers could not say to major customers that the network is being replaced with IP only. IADs are a solution to integrating the communications equipment already installed at the site with a view of migration to IP using intelligent, yet seamless, protocol conversion. This allows rapid roll out of new networks with minimal customer resistance and future services deployment without using a technician.

IADs and intelligent protocol conversion cost-effectively entice end users to migrate to modern IP networks. IADs allow end users to connect their existing equipment onto the modern networks without having to purchase a new office full of equipment.

IADs can offer a fully integrated solution for a business's entire communications network into a single carrier line, significantly reducing communications costs. SMEs today use a variety of separate lines dedicated for specific functions—for example, several public switched telephone network (PSTN) lines for the phone, fax, and personal computer (PC) with Internet connection, a dedicated line for electronic point-of-service (POS) financial transactions, and a leased line for security monitoring equipment. Four or more separate lines dedicated to individual functions. Some IADs on the market overcome this by offering a fully integrated solution for a business's entire communications network. IADs can connect any asynchronous or synchronous devices, such as PCs, with Internet connection, a standard telephone/fax system, security monitoring equipment, POS terminals, and other services to a single network connection (see *Figures 1* and *2*).

FIGURE 1

Typical Australian Retail Infrastructure Set-Up

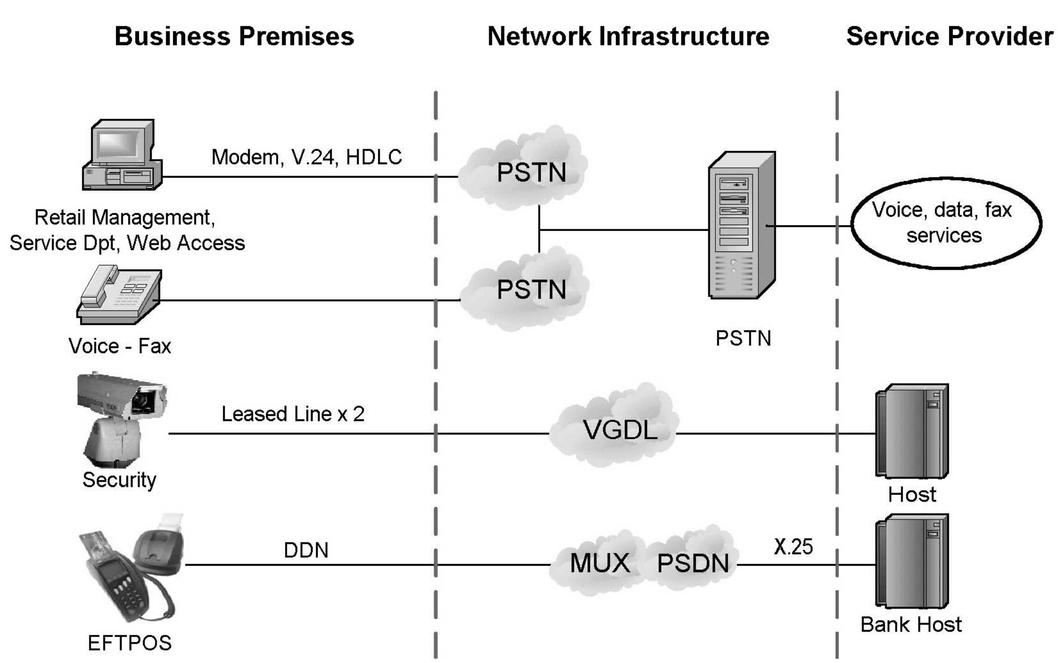

FIGURE 2

IAD Capabilities Will Differ According to Individual Product Specifications

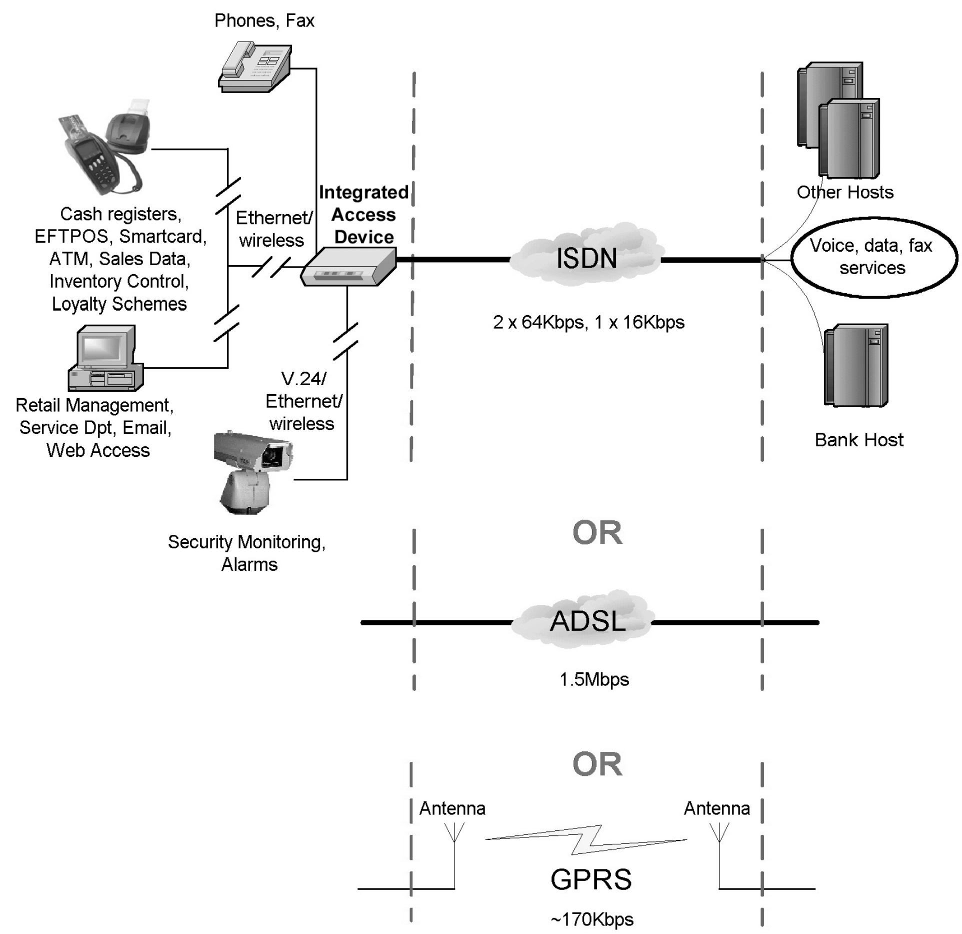

IADs currently on the market are a simple plug-and-play technology that makes installing a modern IP network easier than a new VCR. The plug-and-play technology allows users to plug into the network, immediately minimizing down time. Quick and easy installation and use is another important aspect in migrating end users to modern networks.

The carriers and service providers could offer these access devices as inclusive with the installation costs. Carriers and service providers have the business case to carry this initial cost that will, in the long term, reduce their costs significantly by phasing out legacy networks faster. This will allow SMEs to quickly migrate to modern IP networks that will benefit all stakeholders.

Australia's largest telecommunications carrier successfully addressed these issues by offering an IAD with their integrated services digital network (ISDN) as part of the installation cost. This product enables end users to immediately connect their existing legacy equipment to an ISDN service without the capital cost increase to burden the SME. This solution is currently being successfully deployed throughout the Australian retail industry.

The carrier and service provider can remotely manage and maintain the network from a central management center through the IADs. This enables remote downloading of new software, which allows the simple deployment of network upgrades. This provides an extended life cycle of the IAD by adapting to the requirements of the users and their businesses. Network diagnostics can be managed without having to send a technician to the user's premises. This will decrease the downtime that a user may experience in a network failure.

Network management becomes easier, more cost-effective, quicker in downloading new services, and easier and quicker to maintain—all without the need for field technicians to go to the user premises.

Put simply, the IADs remote management capabilities simplify and centralizes the manageability of an entire network, easily reducing the costs incorporated in managing and maintaining the network. Remote management also provides the ability to move with new directions in communications, significantly increasing the devices potential and life cycle.

A retail petrol seller in Australia usually contains several businesses: gas pumps, mechanics, food stores, and takeout, such as Burger King. With the IAD, it is possible for all of these businesses to use the same connection, significantly reducing communications costs.

These devices will enable carriers and service providers to quickly phase out expensive legacy networks that are fast becoming obsolete. Maintenance and management can now be performed remotely, reducing costs significantly. Both the service providers and the carriers can offer new services to customers that can be offered only on the modern IP networks. The end user can now access the benefits available on modern IP networks in a cost-effective and timely manner.

With the reduction in new communications equipment purchases experienced in 2001, it is important to have the mechanisms to allow companies to access new technologies without the huge capital expense. With this in mind, the importance of protocol conversion and IADs will increasingly play an important role not only for SMEs, but also for the business population wanting to keep up with cutting-edge technology infrastructure.

SMEs want reliability of service, speed, security, simplicity, and upgradeability. Having the latest technology would be great, but on a cost basis it is not feasible for the majority of the SME market. IADs bridge this gap between the modern networks and the end user. All stakeholders in the communications process can benefit from this intelligent protocol conversion device.

Seeing the Real Telecom Future

David Heard
President and Chief Executive Officer
Santera Systems

Few industries have experienced a more jarring change than that which has shaken the telecommunications industry in recent years. But significant change continues in our industry. To see more clearly the future of the telecom infrastructure business, we must first examine the forces that are driving these trends.

Emerging carriers such as competitive local-exchange carriers (CLECs), integrated communications providers (ICPs), and data local-exchange carriers (DLECs) face potent new cost and profitability pressures. They must deal with significant and growing pressure to deliver increased capacity, maintain voice and data revenues, and roll out popular new services, all while controlling both capital and operating expenditures. Even though data traffic is growing rapidly, most networks still rely on voice to produce their largest and most consistent revenue stream.

These drivers are growing at a critical juncture in the history of telco network infrastructures. For many operators, their installed switching technology is at, or nearing, the endpoint of its useful lifespan. For the industry as a whole, the time is rapidly approaching for a major reassessment of its fundamental approach to the telecom switch.

Technology Cycles

As with many other developing technologies, the telecommunications switch undergoes a periodic and largely predictable cycle of evolutionary innovation.

With the advent in 1897 of step-by-step switching systems, telephone companies introduced rotary dialing and enhanced both the privacy and capacity of their networks. A generation later, in 1938, the emergence of crossbar systems reduced both the power and maintenance requirements of the telecom switch. In the early 1960s, the deployment of electronic systems heralded the introduction of transistors, stored memory and speed dialing, call forwarding, and other features.

Then, in 1976, with the roll out of first-generation digital technologies, the industry entered what once was called the "modern" era of telecom switching. Digital switching dramatically increased the carrying capacity of the telecom network, while supporting the introduction of time division multiplexing (TDM) and the eventual deployment of the integrated services digital network (ISDN), advanced intelligent network (AIN), and other enhanced features.

Those current-generation switches, which form the installed base for telecommunications networks the world over, are rapidly nearing the upper limits of their traffic capacity as well as feature and service-delivery capabilities. Astute managers recognize the need for new solutions. The questions now are how to plan for, select, and deploy the coming generation of telecom switching technologies.

Time for a Switch

As is true at every cyclical turning point, today's emerging switch technologies must meet the immediate requirements of current telco networks *and* the future demands of a changing marketplace. Some have described this as the need to satisfy both now-generation and next-generation (NextGen) objectives.

To ensure a smooth migration from current-generation technology to NextGen switching efficiencies, network operators need a switch that delivers comprehensive legacy capabilities and the ability to satisfy a demanding set of emerging performance, regulatory, and survivability demands. That's no small order.

But given the growing competitive pressures in most markets—and the fact that many service providers simply will not install another outmoded legacy switch in their networks—the time is ripe for the emergence of an entirely new type of switch.

With the advent of the softswitch model, network operators will enjoy dramatically reduced costs, expanded density, and the efficiencies of the distributed call-processing model. NextGen switches will be smaller, simpler, easier to maintain, and capable of supporting a still emerging new wave of revenue-generating services.

By providing a single integrated switching platform, this new class of switch provides a scalable, long-term solution for both incumbents and the new competitive carriers. These more advanced switches will combine the core and edge functionalities of both the voice and data networks, and will allow services to be switched in TDM, Internet protocol (IP), asynchronous transfer mode (ATM), or frame-relay formats.

Capable of being deployed at the network edge or at the central office (CO), co-location, CLEC point of presence

(POP), or remote switching center, these more advanced switches deliver superior density and flexibility, and more powerful network controls. This new species of switch will, in effect, bridge the chasm of the looming circuit-to-packet evolution, while opening a seamless pathway to new applications and innovative business models.

A Switch Checklist

Switching is a significant challenge for telecom equipment providers.

Big name producers of traditional switch solutions—hit by layoffs, stock slides, and acquisition rumors resulting from the general slowdown in telecom equipment purchases—seem to have lost the focused drive needed to develop and deploy a new generation of network switch technology. Legacy players have tried to bolt add-on features to existing equipment. But those point-solutions often create little more than stop-gap solutions and new potential network points-of-failure. Many suppliers, vulnerable to shifts in funding and technologies, lack the resources needed to launch the new species of switch technology that network operators now require.

To appreciate the difficulties faced by the supplier community, we should first understand the requirements of a true NextGen switching solution.

Any advanced switching solution must deliver a comprehensive set of Class-5 capabilities as well as a flexible platform for providing enhanced revenue-generation data services. This technology will ideally incorporate both current-generation TDM and emerging-generation packet fabrics, an innovative architecture that will support both native and any-to-any switching modes with unmatched density and growth capabilities.

Specific NextGen switch requirements include the following:

- ***CALEA:*** Under the Communications Assistance for Law Enforcement Act (CALEA), telecom carriers must adapt their networks and services to provide authorities access to voice content and call data communications originating from, terminated at, or forwarded by any subscriber under authorized surveillance. All carriers must meet this requirement, but not all switches provide a simple and efficient CALEA solution.

  While some suppliers address CALEA with expensive and failure-prone network add-ons, true NextGen switches will support a seamless, powerful CALEA capability in both the TDM and packet fabrics. These advanced switching technologies will leverage today's more capable protocols and use media gateways to establish the one-way connections used to implement CALEA bearer monitoring. These new switches will also allow data calls to be tapped and transported for off-board analysis.

- ***Survivability:*** Network reliability and availability will be key competitive factors in tomorrow's telecom environment. To satisfy the survivability demands of the PSTN, this new generation of switches must deliver proven five-nines reliability through the use of no-single-point-of-failure redundancy at the hardware, software, control, and bearer traffic levels.

  Specific reliability features should include physical and logical redundancies, continuous self-test, and sanity monitoring. Key common control hardware and all external fast Ethernet connections must feature redundant, constantly available configurations.

  Feature exchange platforms should also be designed for optimum availability and redundancy protection. This level of survivability is assured through the use of multilayered software and process separation, built-in hot-swappable power converters, and comprehensive, automated monitoring and alarm systems.

- ***Network-Level Survival:*** As subscribers demand ever-higher levels of availability—and considering the threat conditions in the post-September 11 environment—network operators must now also address a far broader and more demanding set of network survivability issues. The basic concept of building separate, virtually managed core switching elements is reasonably well understood, but in today's more security conscious landscape, network managers want switches that offer a higher level of survivability.

  Forward-looking switch designers are now refining a virtual-cluster approach that would deploy two or more mirror-imaged master switch controllers (MSCs) capable of surviving a single disastrous network event. Using active standby coordination, central control functions would be separated into multiple physical locations, each of which is capable of controlling the entire switch fabric. Should an event result in the loss of any MSC, key switching functions would instantly and automatically be switched over to another operational machine.

- ***E-911:*** Carriers must now also deploy switching platforms capable of meeting today's more demanding enhanced 911 service requirements. The best new switching solutions allow carriers to route emergency calls by priority, to bypass network traffic controls, and to quickly pass control of the call from the network to the appropriate public-safety agency. Switches that employ a TDM matrix can measurably reduce the risk, complexity, and cost of delivering enhanced emergency service capabilities.

- ***LNP:*** Originally mandated in the Telecommunications Act of 1996 to enable local competition, local number portability (LNP) is a critical feature expected in today's telecom landscape. Carriers of any size require LNP capabilities to efficiently implement Federal Communications Commission (FCC)–mandated LNP services, so they can manage any magnitude of telephone-number porting orders and changes quickly and reliably.

- ***VoB:*** Voice over broadband (VoB) is poised to drive what many expect will be the next major expansion of service capabilities, and service providers will need switching technology capable of supporting full VoB functionality. The key features needed in a

VoB–enabled switch include ATM access with dynamic routing and high port density; a voice gateway with enhanced density, compression support, and ATM adaptation layer 2 (AAL2) over subscription; and full Class-5 capabilities, including automatic messaging account (AMA) billing and full feature support.

- ***Legacy Support:*** This new class of switch must also provide full support for a wide range of legacy applications, including primary rate interface (PRI) and signaling system 7 (SS7)–International Mobile Telecommunications (IMT) interfaces, transactional capabilities applications parts (TCAPs) for LNP, and simplified message desk interface (SMDI) for voice-mail services.

Coming Dislocation

What kind of supplier is capable of delivering this new class of switch? There are currently four key tiers in the switch infrastructure supply marketplace. Tier One is comprised of traditional high-end players such as Nortel Networks, Lucent Technologies, Alcatel, and Cisco Systems. Tier Two is defined by somewhat smaller switch industry competitors, such as Ericsson, Siemens, and others. Tier Three is known to some as the "Emerging Giant" sector and is distinguished by development-oriented companies that hold the promise of moving upwards toward Tier Two or Tier One status. Tier Four players are traditional start-ups that continually struggle to attract capital but that may someday emerge with viable technologies and business models.

As anyone conversant with today's telecom industry knows, players in Tiers One and Two as described here are currently under tremendous marketplace pressures, with one or more Tier One companies currently divesting themselves of switch-related business units, and Tier Two companies also showing signs of abandoning the switching field to pursue other business strategies. Some have speculated that, within the switching sector, Tier Two will collapse altogether, while Tier One will undergo a massive shakeout that will leave just one or two current players still in place.

For this reason, a number of informed industry observers now consider Tier Three's Emerging Giants to hold the most promise as a source of telecom switching innovation and solutions.

The distinguishing characteristics of an Emerging Giant are best-in-class switch technology and a fully funded business plan. These future winners will also posses unmatched experience in carrier networking and clear operational plans for the deployment of NextGen switch technologies. And, of course, a qualified supplier will offer a comprehensive program to help network operators and service providers understand, integrate, and operate a productive NextGen switching network.

Carriers have a great stake in this coming dislocation, because their networks and ultimately their core business models depend heavily on the availability of workable and competitive switching solutions.

There is little doubt that NextGen switches must and will emerge from the current market upheaval. When they do, they will bring a wealth of benefits to both network operators and their subscribers.

NextGen Advantages

Virtually any incumbent or competitive telecom network can benefit from the qualities of a NextGen switching solution. By integrating easily with legacy systems, interfaces, and applications, these new technologies actually extend the useful lifespan of installed infrastructures. But because they also support a seamless circuit-to-packet migration, these new switches also encourage the integration of third-party services and the rapid deployment of popular and profitable new service features.

As true carrier-class technologies, this new class of switch can deliver the 99.999 percent reliability needed in today's public switched telephone network (PSTN) environment. These platforms can scale quickly and economically to serve as few as 2,000 to more than three million subscribers from a single network cluster. When compared to current or proposed alternatives, the best of these new switches can reduce capital costs by more than 50 percent, while reducing network operating costs by as much as 70 percent.

Conclusion

Those are the driving forces, technology requirements, and promised benefits behind the migration from current- to coming-generation telecommunications switch solutions. Depending on their business model, network operators may be at, or nearing, the decision point that will determine their own switching strategies.

By understanding the issues relating to switching infrastructure, telecom managers can see more clearly the future of their own evolving network.

Data Services Profitability: Learning from the Voice-Feature Business Model

John Holobinko
President and Chief Executive Officer
Aplion Networks

The communications industry is facing a change of epic proportions. The underlying revenues for basic voice and data services upon which this industry was built are under a virtual siege. And there are no killer applications identified on the horizon to bring the industry out of its current malaise.

Beware of What You Wish for…

The greatest service providers' dream was creating cheap bandwidth. Yet, cheap bandwidth has provided none of the economic benefits that the industry anticipated. Instead, cheap bandwidth has led to oversupply and rapidly falling prices. Profit margins have not increased. Instead, they have been declining. Prices for bandwidth no longer reflect the cost of providing that bandwidth, but instead only what the customer is willing to pay. Enterprise customers embraced the concept of cheap bandwidth, not to buy more of it, but rather because they view the cost of bandwidth as a necessary business overhead whose cost should be reduced as much as possible.

Technology Advancements Are No Longer an Advantage

For the first time, technology is changing faster than the industry's ability to absorb these changes. The ability to expand transmission bandwidth capacity is no longer limited by technology, but rather by today's economic conditions. Those who have argued that the industry's economic bottleneck is local access must explain why service-provider local access networks can be purchased at pennies on the dollar (e.g., Winstar's assets) as service providers go out of business.

Convergence has been an utter economic failure. While service providers and venture capitalists embraced the idea and built solutions to transport data, voice, and even video traffic on the same wire using Internet protocol (IP), enterprise customers have overwhelmingly rejected the business proposition. Proof is that today, frame relay and asynchronous transfer mode (ATM) services revenues continue to grow faster than IP services revenues. In the eyes of business customers, the loss in security and performance by leaving these trusted services was not worth the benefits that IP transport promised.

The Industry Challenge: Incremental Revenue, Profitable Services

If a reversal of the current economic malaise in the industry is to occur, it cannot be through restoration of the old basic voice/data bandwidth services model. There is no single killer application that will save the industry. If there were, it would already have emerged in eight years of industry searching since the Telecommunications Act of 1996 created the basis of today's market.

Rather, the challenge of the industry is to gracefully evolve to a new economic model in which communications services are easy to offer and maintain and have high profit margins. Yet what kinds of services offer this potential and what types of customers are willing to buy these services?

Data services represent more than 80 percent of today's consumption of bandwidth, yet provide less than 20 percent of overall service revenues on an industry basis. On a profitability scale, the numbers are even worse. Clearly, data services represent the greatest opportunity for increasing revenues and profits (see *Figure 1*). Correspondingly, enterprise customers have the most to gain from communications services and are therefore willing to pay more for data services if they provide more direct business value. This is a big "if." Few of today's communications services provide direct customer value. To provide direct value, the service must do one of three things: It must improve enterprise business efficiency; help them serve their customers better; or provide functionality that they cannot otherwise afford.

Spiraling IT Costs: The Industry's Best Friend

One of the fastest growing areas of enterprise costs is for information technology (IT) services (see *Figure 2*). Even as

FIGURE 1

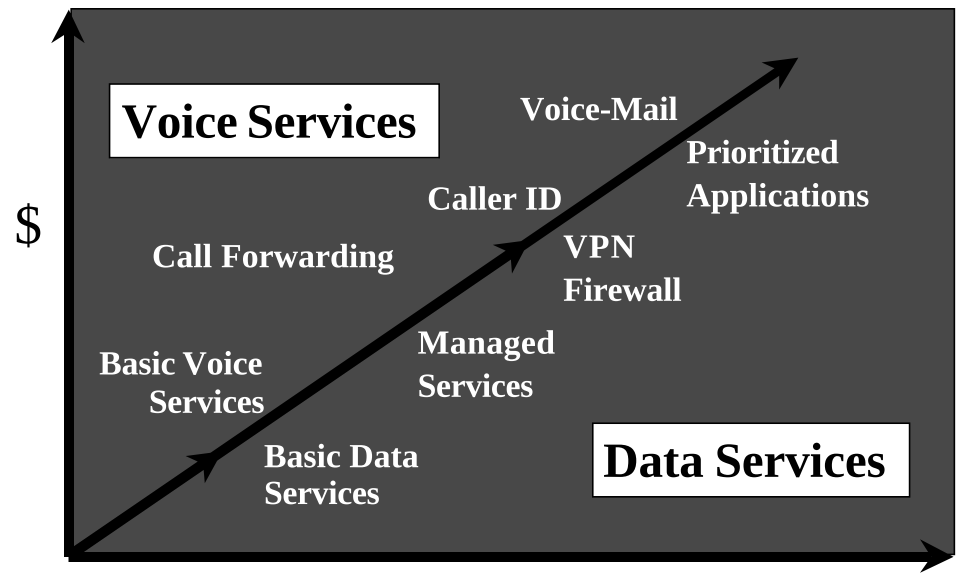

Managed data/ packet servi ces need to be as pro fitable as managed voic e features.

the economy has weakened, enterprise IT expenditures are not falling. IT professionals are expensive. New IT applications have high up-front costs. It is this condition that represents a large business opportunity for services providers: The ability to provide managed services to enterprise customers represents a potentially huge area for revenue growth for service providers, because it displaces more costly IT expenditures for the enterprise customer. Increased managed services revenues to the service provider represents either a cost savings or greater return on investment (ROI) to the enterprise customer, compared to the expenditures for hardware, software, and IT support personnel that they would otherwise have to make.

Managed Services: Great Potential, No Profits

Many service providers have known that managed services are a market that they must address. Yet, they have had lit-

FIGURE 2

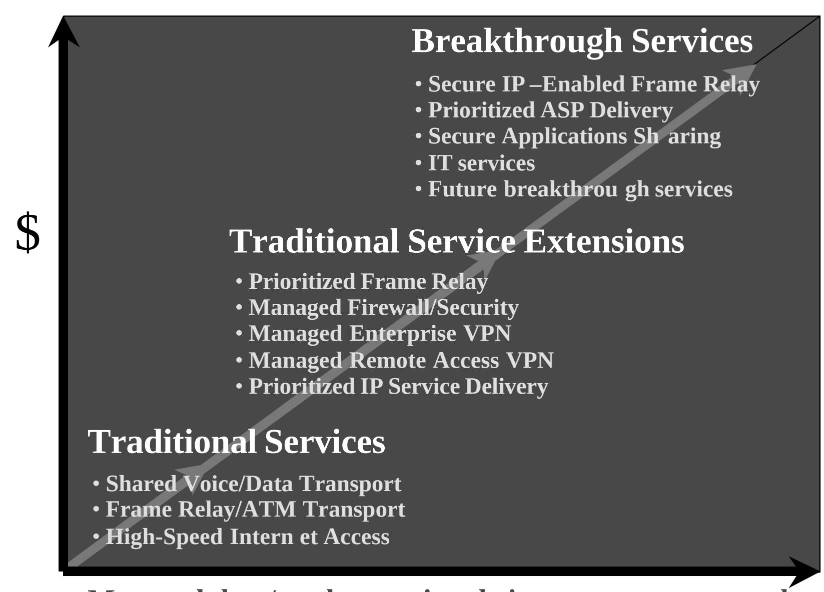

Managed data/ packet servi ces bring gr eat revenue a nd profit opportuni ties.

tle, if any, success in providing these services. Historically, they have been costly to implement and difficult to maintain. As a result, the prices that must be charged to customers make them unattractive to the majority. And for many of these services, the ability to guarantee service levels on a continuous basis has been difficult or impossible. Therefore, managed data services have continued to be a source of service-provider difficulties.

It is easy to understand why the underlying cost to provide these services is so expensive. It is not the network infrastructure, which represents only 20 percent of service costs. The life-cycle costs to support these services represent 80 percent of their cost. A managed service typically requires reprovisioning of the network; additional equipment at the customer premises, central office (CO), or both; and expensive truck rolls in order to implement and maintain each service (see *Figure 3*). With these high up-front and sustaining costs, most managed data services have only a marginal chance of being profitable, and then only in out years. Service contracts must be typically two years or more, which further reduces customer acceptance.

Some service providers have outsourced selected managed services to companies that specialize in providing one or two managed services, such as IP–based virtual private networks (VPNs) and firewalls. In this scenario, the outsourcer installs and manages the service. The service provider has no control over service costs, and its revenue consists of only a small percentage (e.g., 25 percent) of what the customer pays for the service. This limits the service provider's ability to generate additional profit margin by leveraging pricing, customer penetration, or underlying service costs, because they are all outside of the service provider's control. Furthermore, responsibility for service problems is still placed squarely on the service provider by the customer.

Data Services Management Is Too Complex

The equipment in today's networks, whether connection-based or packet-based, does not have the ability to manage actual services. Today's networks—including the switches, routers, and multiplexer equipment—cannot themselves recognize the difference between applications that run over bandwidth. They can only manage and prioritize traffic as defined by circuits, not the actual traffic itself. Therefore, to provide quality of service (QoS) requires a separate circuit to be managed across the entire network. Whether these circuits are frame relay permanent virtual circuits (PVCs), ATM virtual circuits, or IP–based multiprotocol label switching (MPLS) tunnels, the complexity and the overhead are similar, independent of technology.

Recognizing and prioritizing by application takes additional equipment for each customer. For the service provider, installing, integrating, and maintaining this equipment for each managed application adds cost and complexity. The burden for integrating multiple hardware and software elements necessary to instantiate the service rests on the service provider. When multiple data services are implemented, features of one service often cause unintended interactions with another of the services, resulting in service problems. It is no wonder that service providers find that the initial cost of offering these services makes them initially unprofitable. As the number of customers rises, the service life-cycle costs do not scale, but rather, in many cases, get worse (see *Figure 4*).

The "Best of Breed" Fallacy

Historically, service providers have evolved to an equipment selection process in which each network element (NE) in the transmission network can be selected independently of other NEs. This gives the service provider maximum flex-

FIGURE 3

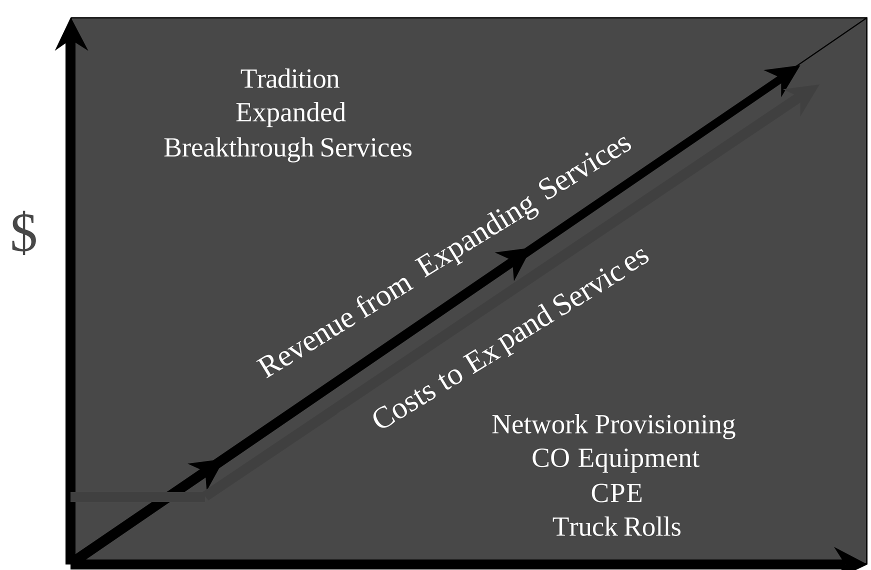

The life-cycle costs of managed services are obstacles to success.

FIGURE 4

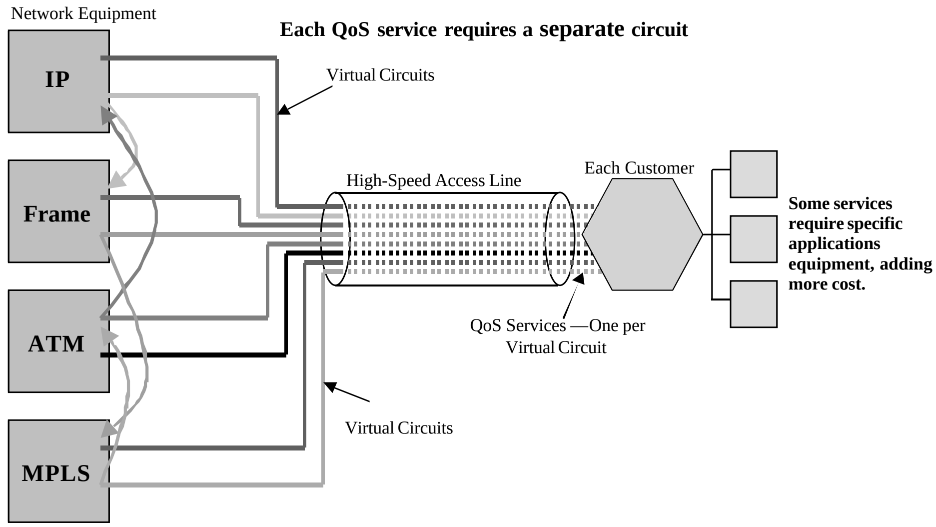

Why are current service costs so high?
Life-cycle costs are 80% of the total cost.

ibility. This "best-of-breed" selection process works for traditional bandwidth-based services because transmission-layer technology is very mature, including standards that have been developed over many years.

However, managed services are relatively new. In the case of managed data services and advanced applications services, there exists no mature model upon which all equipment and software can conform to at the *service level*. Therefore, the service provider is responsible for all service integration and the resolution of conflicts due to interactions between various vendors' software and equipment (see *Figure 5*). In the case of advanced data services, "best of breed" has not been a successful business model. It has resulted in the high service life-cycle costs that prevent these services from being priced reasonably and becoming highly profitable.

A Proven Model for Success: Voice Features

Voice features continue to provide the highest profits of any communications services. There is a good reason for this. When implementing these services, the service provider

FIGURE 5

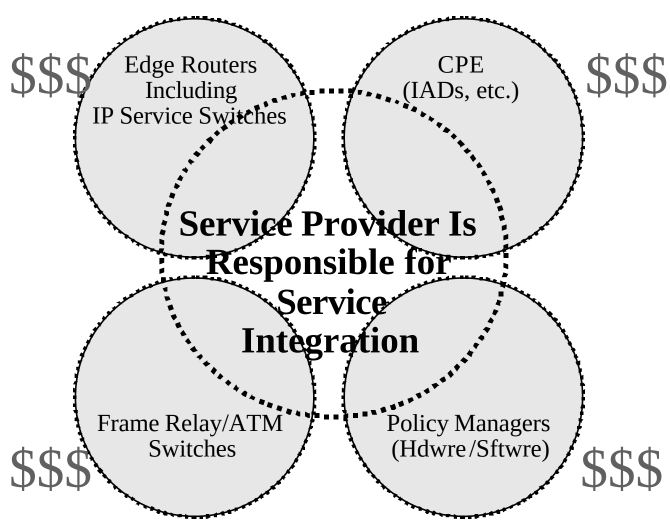

Equipment integration is complex and expensive when deploying managed data and packet services.

bears little of the costs associated with installing and maintaining advanced data services. For example, when a customer orders call forwarding, the service provider does not have to provision a new circuit across any part of the network. The service provider does not have to install new equipment. A service call does not have to be made to the customer—no truck roll is initiated. Instead, these services are installed, activated, and maintained simply by a point, click, and deliver software process.

The reasons for the unequalled success of voice services are important to the data industry. When advanced voice features were initially introduced, individual voice switch vendors provided a complete solution at the service level. The vendors, not the service provider, took on the total responsibility to integrate the necessary software and hardware to implement these services. The equipment vendors also provided full interoperability between the various voice service features that they offered to insure that changing a parameter in one service did not deleteriously affect another service. Thus, single vendor solutions enabled the service provider to implement services rapidly and inexpensively and to generate profits that continue to today.

A Managed Data-Services Model Based on Voice Features

There is no question that customer demand exists for managed data services, given that they can be priced attractively compared to enterprise customer IT costs. The challenge is to remove the current costs and complexities that make them so pricey and difficult to manage.

The goal is to have the capability to implement managed data and packet services with a simple Point-click-*Deliver*[SM] model, similar to how voice-feature services are implemented (see *Figure 6*). For managed data services, Point-click-*Deliver* means eliminating circuit provisioning, additional equipment, and truck rolls each time a service is sold and delivered. In other words, generate highly profitable service revenues now.

To accomplish this requires that a single vendor take on the responsibilities for providing a services solution that eliminates specialized equipment for each service. It requires a solution that can enable the current network to become a smart, software-enabled network that can recognize and prioritize *services* instead of simply managing *circuits*. The solution must relieve the service provider of the complexities associated with service integration. This can now be accomplished within the service provider's existing network infrastructure, current customer connections, and without affecting existing customer services.

Such a solution consists of three main hardware and software elements (see *Figure 7*):

- An intelligent services switch at the network edge that is interoperable with frame relay, ATM, IP/MPLS, and time division multiplexing (TDM) equipment
- An inexpensive intelligent service agent at the customer premises
- Intelligent service software that runs at the service provider's network operations center (NOC), which reduces the installation and implementation of services to simple software commands that are downloaded to the intelligent service switch and service agent whenever a new service is desired

If all three elements of a smart network are provided by a single equipment manufacturer, then installing and managing services can be made as simple as installing voice features. Customer connections can be provisioned once at the time of initial connection. Service speeds and priorities are then controlled by the simple software commands to the equipment on each edge. Circuit provisioning across the

FIGURE 6

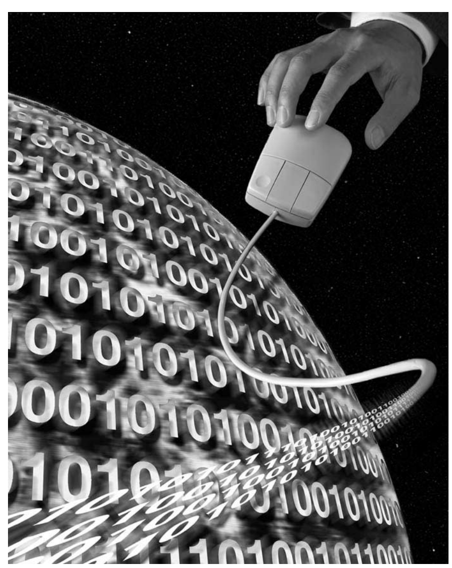

FIGURE 7

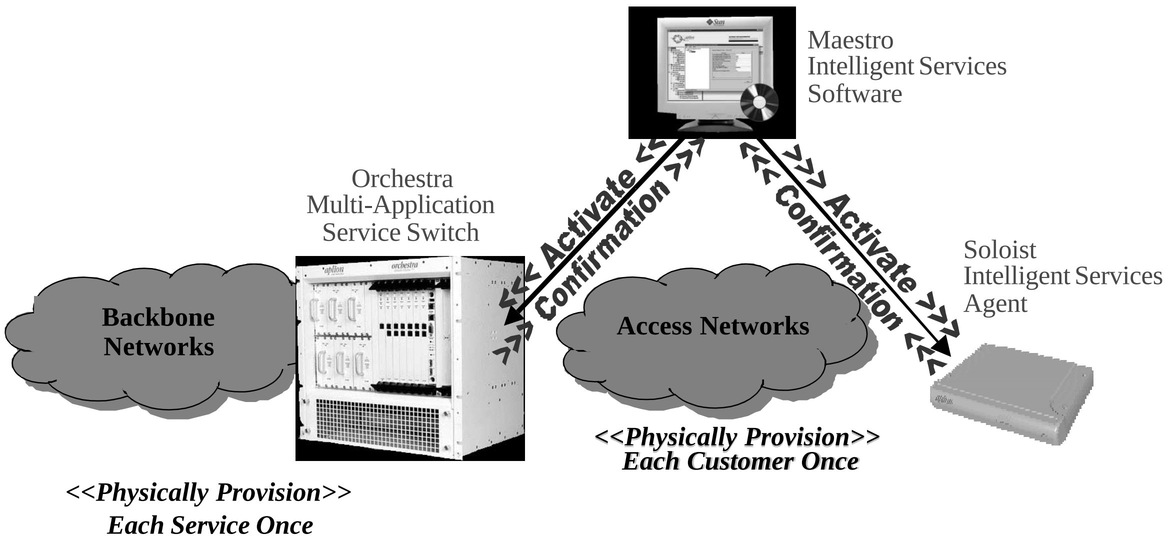

network backbone and through the access network can be decreased by more than 90 percent. Each time a new service is implemented, no additional equipment is necessary. The new functionality is simply downloaded to the equipment. Similarly, changing technology and service definitions can be implemented without the expense of changing hardware.

By eliminating the major service life-cycle costs of managed data services, they can be priced reasonably and generate excellent profits. Managed services can be aggressively marketed with free customer trials, because the cost of installation becomes so low.

Resistance to the Voice-Features Model

There are many service providers that have criticized the voice-features model. They argue that having dependence upon a single vendor was at times frustrating for the service provider. They had to wait for a vendor to complete a feature before it could be offered. It limited their flexibility. These arguments are unquestionably valid. Yet during the last 15 years, voice features have provided exceptional industry profits. It is only after many years of stability that next-generation software-based voice switches have emerged, which enable the service provider more flexibility to control their own services and to combine features from different vendors ("best of breed").

Now compare this to managed data services, which from their infancy have been implemented using a "best-of-breed" model, where individual NEs from different vendors are used to create, deliver, and prioritize a service. For the service provider, integration complexity has skyrocketed. Profits are virtually nonexistent. Services take weeks, if not months, to implement. Initial service installation costs are so high that most services do not break even for more than a year.

Clearly, a "best-of-breed" process is attractive when services are mature and already generating high revenues and large profits. It has not been successful for emerging services, such as managed data services.

It's Time to Learn from Voice-Features Business Success

The success of the voice model is unquestionable, while the "best-of-breed" model for managed data services has been a complete and utter failure. In retrospect, this should not be surprising, given that voice-service features did not move toward a "best-of-breed" model until they were extremely mature. For managed data services, today service providers need a complete managed services solution. They need a Point-click-*Delive*r model that gives them a simple, flexible, and inexpensive means to offer an extremely wide array of managed data services, easily configurable to each customers specific requirements, without network and equipment constraints.

With the rapid decline of basic voice-service and basic data-service revenues and profits due to the "cheap-bandwidth" conundrum, service providers need to quickly enhance their revenues at low capital cost. Managed data services represent a powerful engine for additional revenues over existing networks and customer connections. The service provider's fastest path to new revenues and profits is to move to a smart-network model in which a single vendor provides fully integrated managed data services solutions. A Point-click-*Deliver* process for managed services is much more lucrative for rapidly increasing revenues and profits of data and packet services as well as customer acceptance.

The Effect of Fast Technology Transitions on Companies' Survival

Abdelmaksoud A. Khalifa
Chief Executive Officer
Quttan Electronics

One of the most important things that should worry most chief executive officers (CEOs) now is product life cycles in almost all industries. It has shrunk, but in no industry has it become shorter than in the electronics industry—especially within the communication sector.

Customers are very demanding now and driven more by trends than ever before (thanks to the the information revolution and the Internet). This is creating risks as well as opportunities for companies within the communications field. However, the short-term payback from research and development (R&D) has been disappointing for a very high percentage. An increase in many companies' rate of cash burning, which outpaces their rate of income, is threatening these companies' survival.

The industry's shaking is leaving no one at rest. There are some exceptions in countries where communication companies are well protected by the government, but even this may change as the World Trade Organization (WTO) attempts to enter the doors of all the world's countries within the next few years.

We can take as an example the unrest in the home entertainment industry. Initially, cable television was enjoying stable profitable position, until mid-1980, when new technologies arose. Now there is competition between cable on one side and assymetric digital subscriber line (ADSL) and interactive digital satellite systems on the other side. If we were to compare the infrastructure cost differences between these systems, we would find how difficult it is for the cable industry to manage these costs now—how difficult it is for cable television (CATV) operators that have had to cut their charges in order to face the challenges brought about by the arrival of new technologies. Cable companies are in a bad situation now because of the high amount that they have heretofore invested in infrastructure as well as licenses. If they migrate to the new technology, they would not make use of their current assets. At the moment, they still have good market-share, but the competing technologies are becoming ever cheaper and are slowly but steadily eating away at that share. This paper is not intended to focus on the entertainment business, however this is a good example of how fast development in technology will affect companies' long-term investment and payback.

In the communication industry in general, the picture is more dark than in the cable television and home entertainment industry, because customers' tastes in their means of communication are always quickly changing. Perhaps one of the reasons why the cable television industry is not collapsing may be that end users do not care how their television channels or broadband arrive to their televisions—be it via special cables, air DSTV, or a telephone line—all they really care about is picture quality, price, services, types of programs, etc. But, conversely, the general communications client is directly interested in handset size and fashion, download speeds, mobilization and roaming facilities, etc. This is perhaps why they are quicker to migrate to the newer technologies in communications than customers in any other industry. This fast transition is the main problem.

Companies could face two situations in this transition:

1. ***Soft Transition***: This will mainly depend on the technical differences between the new and old technology. If the difference is not wide and the old technology could be used as base structure for the new one, then the transition could be called "soft" and will not threaten the company's survival or financial health. An example of this would be the cellular-phone change in Europe from the extended total access communications system (ETACS) to the Global System for Mobile Communications (GSM) standard. Despite the fact that the transmitting and switching equipment was totally different, the operator could still use a large part of its assets—mainly its towers,buildings, billing system, etc.
2. ***Hard Transition:*** In this case, the new technology comprises novel technical concepts that do not relate to the old technology, but functions better and is, most of the time, cheaper. One example of this would be voice over Internet protocol (VoIP) versus voice over normal cable or satellite channel for long distance. The old technology is based on 4KHz telephone channels over, for instance, satellite, while the new voice over Internet technology is a totally different system

that depend on gateways and routers to manage calls. Besides there being a huge difference in the investment between the old and new technology, there is also a great difference in that any company from any country could operate the new technology with no limit. We could conjecture that X company, located in tax heaven, could perhaps successfully compete with, for example, AT&T, which may be paying more taxes and licensing fees. Furthermore, labor costs would be different, as well as material costs. Such a dangerous transition could casue companies with the old technology to go bankrupt. The financial statements of AT&T for the last two years is a good indication of how they are feeling this heat—but thanks to their very large market-share and financial health, they can buy time to change.

To face this dilemma, there are several options available to companies. None of them is 100 percent failsafe, but each company should choose the one that best matches their position:

- ***Avoid projects with a long payback period.*** By doing so, companies will minimize the risk of having their investment in a technology become obsolete, and their initial invested capital will be recollected from the project before any unhappy surprises can arise. The disadvantage of this option is that big projects never have a short payback, and choosing only short payback projects could lead a company's market-share to dwindle and the company's size to shrink.
- ***Establish a very strong R&D department.*** The advantage is that this team will lead company creativity, and the company will thus generate new products frequently. Moreover, the R&D department could improve the company's understanding of new technology development and could also work as early warning system for any new technology that might threaten the company's current technology. The disadvantage is that R&D is costly matter and thus only companies with strong financial health can have the lead in R&D.
- ***Focus on services, not production***. The advantage of this is that the company does not have much exposure regarding direct investment in a technology. Because it only does marketing, the company does not have any investment in technology development. As it markets the technology available to it, the company can shift very easily to a new technology. The disadvantage is the company doed not have any competitive advantage, as it is promoting a public technology—competition is very tough, and the company's margin could disappear or become very slim.
- ***Diversify.*** Companies could diversify their business within the communications sector into different sectors in different markets. The advantage is that any loses due to a technology change could be offset by margin from other divisions. Due to their diversification, such companies will deal with variety of clients, which will thus reduce the impact of one division's clients' migration to a new technology. The disadvantage is that such companies could be marginal in all sectors—and not leaders in any of their business portfolios—which will make them weak against the market leader in each segment.
- ***Look for market niche.*** Companies should look for market niche because they can generate good profit from that niche. One example would be the long-range cordless phone, which is a market niche that is served by less than five companies worldwide. These companies control their technology and their customers. The advantage of this that the profit is good and technology change is forced by the companies themselves. They choose the time and capital for the investment in a new technology and are not forced to follow a market trend. The disadvantage is that if growth in a market niche is slow or nonexistant, those companies that are used to having a nice percentage of profit could find that this niche is vanishing or has vanished, thus forcing them into the open market, where they may not survive the mass-market games.

Many CEOs are now focusing more on survival than on making money. A lot of focus is on the stock market and making the stock-market analyst happy, as they are the key for achieving the needed capital for investment in this year's new technology. This is fine for now, but what about next year's new technology. Companies cannot continue investing in new technologies indefinitely—they need time to rest, to get paid back for their investment. This will never happen with companies' merging and becoming bigger and bigger—which provides short-term relief through cost-cutting due to the merger of manpower, but in the long term becomes much more complicated.

One solution might be a collaboration of R&D involving all leading companies in the field. In this way, they could enforce the latest technology and the huge R&D cost could be shared among the group members. One example of this is the development of the Bluetooth technology, which five big companies have been leading. Unfortunately, the launch of Bluetooth has not gone as expected, because costs were five times higher than were hoped. Nevertheless, one collaborative group's troubles does not guarantee trouble for all who endeavor to collaborate in such a manner.

DSL Deployment: The ISP Perspective

Lorraine Lopez
Vice President, Broadband Operations
EarthLink

This paper presents a case study of one Internet service provider's (ISP's) experience with digital subscriber line (DSL) deployment. The ISP serving as the vehicle of the study is EarthLink, currently the second largest ISP in the United States.

Key Strategies

ISPs are vital to broadband's success, but ISP customers are not insistent on a particular kind of Internet access. The customers demand fast Internet access, but if DSL is not an option in a given area, the ISP's proposal of access through cable or a microwave multipoint distribution system (MMDS) will be readily accepted.

Ongoing communication between the ISP and its customer is essential to maintain the ISP's reputation and market share. The quality of a customer's experience should be at the top of the ISP's list of strategies, followed closely by the aggressive expansion into new DSL markets and the concurrent pursuit of cable network agreements.

EarthLink focuses on services rather than on carriers and uses multiple vendors and technologies to achieve its ends. In this way, it provides its customers with access to the Internet and with important services such as e-mail. In addition, the type of access can be changed with no impact on the customer.

The services offered by EarthLink are competitive with those offered by other ISPs: high-speed asymmetric DSL (ADSL) or symmetric DSL (SDSL) access to the Internet, self-installation, dynamic or static Internet protocol (IP) addresses, dial-up access, e-mail, and 24-hour technical support. To attract new customers, the ISP offers free equipment, set-up, and installation in the United States. Unlike some ISPs, EarthLink has found that its customers are sensitive to pricing and can be lured away by attractive offers from other companies. The offer of free equipment and installation combines with the ISP's pricing strategy to keep its customers from straying.

Variable Degrees of Success

The relationship between ISPs and local-exchange carriers (LECs) has often been tense, as LECs have traditionally regarded ISPs as competition. Gradually, however, LECs are recognizing that ISPs are their best customers, and with these improved relationships market expansion and penetration go smoothly.

Although automation of the ordering process is improving, significant problems remain. Different systems are not necessarily compatible, and if a vendor's ordering system has a system enhancement that does not match the ISP's, even a minor glitch can wreak havoc. All parts of the chain—the ISPs, the LECs, and the incumbent LECs (ILECs)—are affected by system changes in the other parts, and those incompatibilities can significantly slow down business.

Technicians are gradually coming to understand their installations more completely. A few years ago, technicians were oriented specifically toward either hardware or software, and the person who physically installed a mast and then climbed the mast to install a parabolic dish was not the same person who entered the house, loaded the software, went through a checklist, and talked with the customer. Today, a single person frequently performs that wide range of tasks. Extensive training programs have helped technicians to be increasingly competent and comfortable in both worlds.

Self-installation is making inroads in the industry, as it eliminates an important number of truck rolls. With the proper guidance and support, self-installation can be very successful, and the costs of deployment will consequently be cut; however, a number of truck rolls will always be necessary. Some customers would rather pay for a truck roll than spend the time involved in self-installation. Even so, the savings in truck rolls and technician time offered by self-installation are significant: The average customer call for self-installation lasts only 10 minutes.

In fulfillment, too, ISPs and vendors often fall short, as only a few people in the industry today have expertise in this area. EarthLink has fewer problems in this area than some, as 98 percent of its customers in Genuity territory are doing self-installation.

Exception handling and escalations are two familiar elements in any ISP's business. The unusual position of the ISP between the LEC and the customer often requires the ISP to approach the LEC with a request to expedite a customer's account to deliver service as quickly as possible.

An ISP's Deployment Concerns

The biggest challenge facing ISPs is living up to market expectations. Customers expect DSL to be like narrowband—easy to install and inexpensive—and that is not currently the case. DSL service is plagued by communication problems and ongoing system troubles punctuated by frequent upgrades and software "improvements."

The customer's experience is further complicated by potential finger pointing between LECs and ISPs. Most customers care less about a problem's origin than about its solution. Customers are often faced with what appears to be unresponsiveness on the part of their providers, and that leads to frustration. DSL can appear extremely complicated from the customer's point of view; in contrast to the troubles associated with DSL service, the customer can call a cable company and be installed in only a few days.

To remedy the difficulties of deployment, ISPs would like to see better coordination between industry partners coupled with on-line sales, provisioning, and order tracking. Ideally, a customer should be able to track his order in detail—for example, logging on and seeing that the service has been ordered, but that the line presents problems and is being checked, or learning that the equipment has been shipped and that the filter and software will follow soon.

In addition, ISPs would like to see loop turn-ups in 24 to 48 hours, and they need enough equipment and capacity to handle delivery of that service.

Generally, ISPs find that DSL is a good service, but customers want fast access and do not care if it comes as DSL or cable. To keep customers, ISPs must keep them informed of their order's progress all the way from the original request for service to the point at which the service is up and billable.

IP over Optics: A New Approach to IP Service Transport

Yaki Luzon
Founder and Vice President, Marketing
PacketLight Networks

Michael Mesh
Founder and Chief Technology Officer
PacketLight Networks

Internet protocol (IP) traffic uses heterogeneous networks and equipment and runs over multiple topologies. The goal of network designers is to handle IP traffic in light of changes and dynamics in networks and technologies in the marketplace.

This paper describes an approach to the transport of different types of IP traffic. The focus of this approach is the metro transport, though it can be expanded beyond that.

Challenges in Transferring IP Traffic

IP traffic has many faces (see *Figure 1*). It can come as a T3 or Gigabit Ethernet (GbE). It can go into the network as a packet of synchronous optical network (SONET) or asynchronous transfer mode (ATM). In many cases it is transported over a SONET network, or SONET rings, without being connected to the existing metro network that carries other services such as voice.

Legacy Networks

A large amount of legacy equipment exists and must be dealt with, whether it is in the edge or the core of the network. There are also multiple protocols. IP runs over ATM and over SONET. SONET is either channelized or concatenated; it runs T1 or T3 rates and up to optical carrier (OC) signals (OC–3, OC–12 or OC–48). GbE is an emerging protocol in the metro network that is targeted mainly at the transport of IP traffic. All of these have to be accommodated simultaneously.

Typical Transport Solutions Today

In many cases today, the network solution is based on converting the different protocols into the protocol used in the core of the network, transferring them into SONET. These interfaces are actually converted twice. Usually Layer-3 processing at the edge is done in order to get into the transport network, since it must be accommodated to SONET transport. At the other end, the traffic has to be converted again to connect to the core service networks in the central offices where the core of the IP network is located (the routers).

Processing in these cases may include quality of service (QoS) handling at the edge and protocol conversion, as well as concentration of the IP traffic into the optical network.

Principles Of IP Service Transport

A new set of principles is necessary:

- Transfer IP traffic across the metro network in its native mode (protocol agnostic transparent edge)
- Take multiple sources and aggregate them in a central point, rather than handling them on a one-by-one basis from each node, to save on the connectivity in the central offices
- Process once at the central location, as opposed to at each and every node in the network, to simplify the edge of the network and save cost
- Switch by service type and not by a specific interface that comes in. It is possible to bundle multiple interfaces and switch them according to the type of interface that they were brought into, their source, or their destination. This allows more flexibility in managing the network.
- Assure that any type of service can be mapped to a wavelength, and multiprotocol label switching (MPLS) and generalized MPLS (GMPLS) can be used as a control plane
- Use MPLS tags to mark different types of services, and provide traffic engineering and QoS

The proposed method utilizes the above principles, taking all the different types of interfaces and mapping them in their native mode into the fiber by packetizing them—the transparent-edge concept (see *Figure 2*). It also takes advantage of the high capacity of the fiber by not concentrating traffic before mapping it to the optical network.

Service Transport

Processing is no longer needed on every node, and transport is not limited to any specific topology. Fiber-optic connectivity allows a combination of topologies including rings, trees, and mesh. In many cases, the data networks are mesh in nature because of the equipment that exists for handling

FIGURE 1

Networks: IP Over…

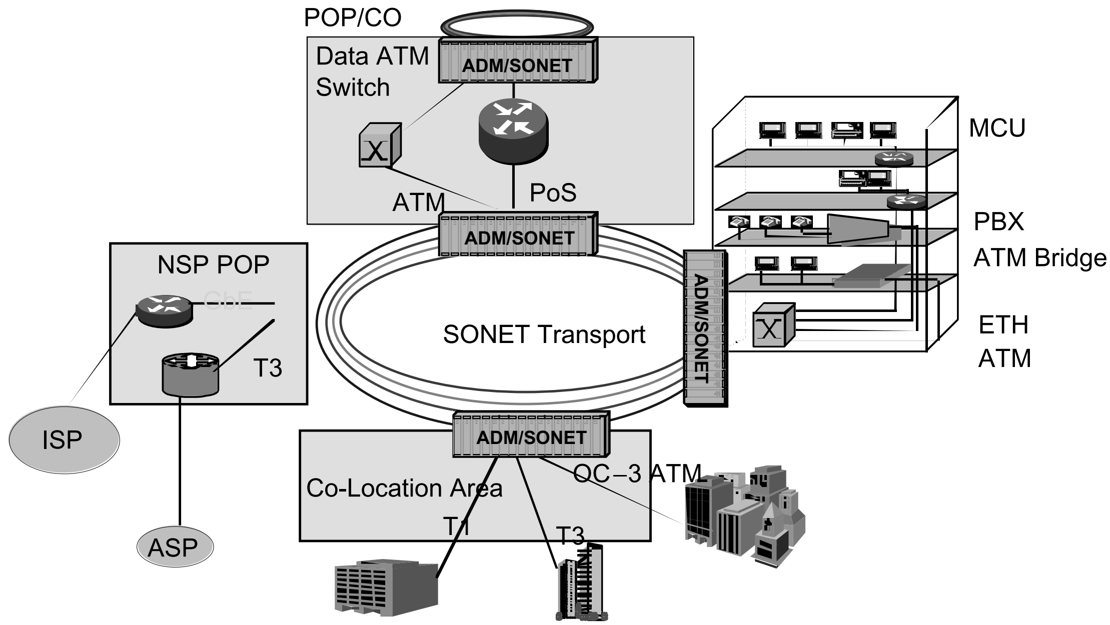

data. On the other hand, the transport network is mostly built as rings. Wavelength connections with dense wavelength division multiplexing (DWDM) can be used to create mesh networks over ring topology. The connections to different types of networks are done at the central location (see *Figure 3*). The IP traffic transported can be transferred over pure IP networks to the routers, all-optical networks to other locations, ATM as a router interface, or even TDM networks that come out of edge routers and are aggregated in the central location.

Using MPLS for Multiservice Transport Solutions

An MPLS control plane is used to connect between service applications and service core networks (see *Figure 4*). This is done in stages. The first stage, at the edge of the network,

FIGURE 2

Transparent-Edge Concept

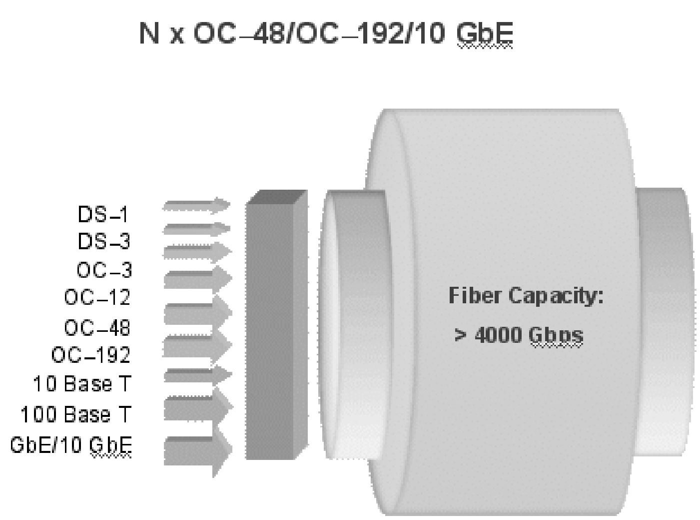

FIGURE 3

Service Transport

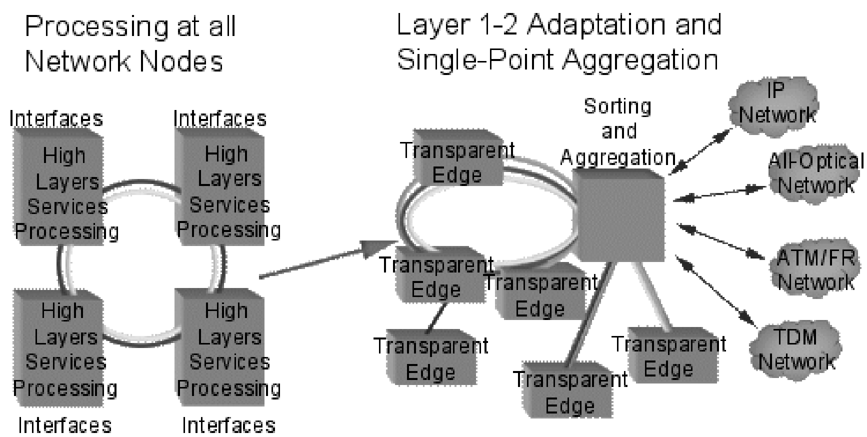

uses multiplexing and MPLS tags to mark the physical-layer protocol and the specific QoS definition; the multiplexed traffic is sorted and mapped based on this label addressing to the different types of network on the other end—the core of the network. Between these two steps, traffic is aggregated and mapped to wavelengths maintaining the MPLS tags. Service switching at the core maps the traffic to the particular network or to other transport networks for remote, central-office connectivity. This service switching uses these MPLS tags as switching control. The end result is that different types of IP traffic collected at the edge can be combined and carried through the metro, then mapped to the particular types of IP networks.

Control Plane MPLS and GMPLS

MPLS and GMPLS provide the capability for traffic engineering and QoS control. This allows for the optimization of routing at lower layers—Layer 1 and 2—without going up into Layer-3 IP traffic and trying to map data from those to MPLS.

This enables the creation of connections based on MPLS protocol. And in the future, when the network becomes all-optical, it will allow the use of the same principle protocol to allow service switching directly over optics and smooth migration to these all-optical networks.

FIGURE 4

Using MPLS for Multiservice Transport Solutions

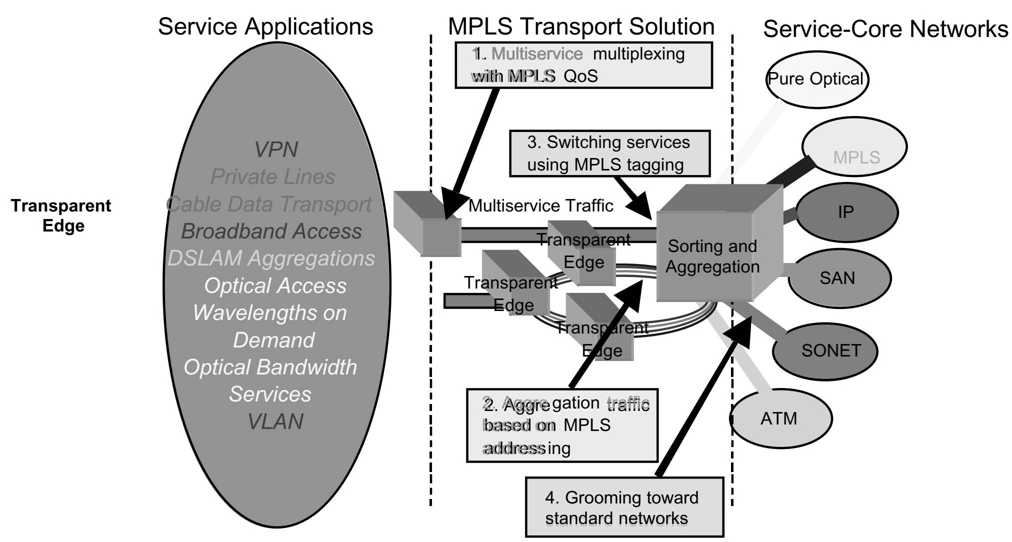

When connecting into the core of the IP network, it is possible to choose the connection by the application and the application type it is coming from. There are multiple methods that correspond to the specific nature of the original traffic. If it is Ethernet or GbE, it will usually be kept at the Layer-2 stream. In some cases of packet over SONET (POS), in order to aggregate these traffic flows they need to be terminated at Layer 3 and then connected to the router at the central office.

In other methods, this can be handled through tunneling; in time division multiplexing (TDM), the channelized TDM is groomed into higher bandwidth channels. All of these are targeted to collect the IP traffic at the edge and connect it through a minimal number of interfaces at the core.

A specific example is that of an Ethernet aggregation (see *Figure 5*). Edge devices that are virtual local-area network (VLAN) aware or unaware can transfer traffic through a transport network, while maintaining the original network application and network structure. A combination of VLAN translation and MPLS tagging allows the transfer of data through the transport network and aggregation on the other side while preserving the original nature of the data flows. Therefore, carriers can have the original structure and type of network application going across the metro network with minimal connectivity at the core of the network.

Similar level of aggregation can be done for TDM traffic.

Summary

This proposed solution is agnostic to the IP physical layer or to the protocol used to transport the IP traffic. It allows the use of different fiber-topologies. The transparent edge provides simplicity and low cost at the edge. It also allows for the aggregation of traffic from different sources in the central office.

The use of the MPLS control plane allows the introduction of GMPLS in order to migrate to an all-optical network and, in the future, run IP over fiber as well as different types of physical-layer services. MPLS provides the ability to handle traffic engineering and QoS, to use tagging across the network, and to transfer the different types of services to the relevant networks. The system is open, since there is no processing at the edge or conversion to a specific protocol. The introduction of new types of transport mechanisms is therefore very easy.

What does this mean to the carrier? It provides simplicity in planning and provisioning at the edge, immediate and future cost savings, and future-proofing by allowing migration from any existing network to an all-optical, all-packet network.

FIGURE 5

Example—Ethernet Aggregation—VLAN

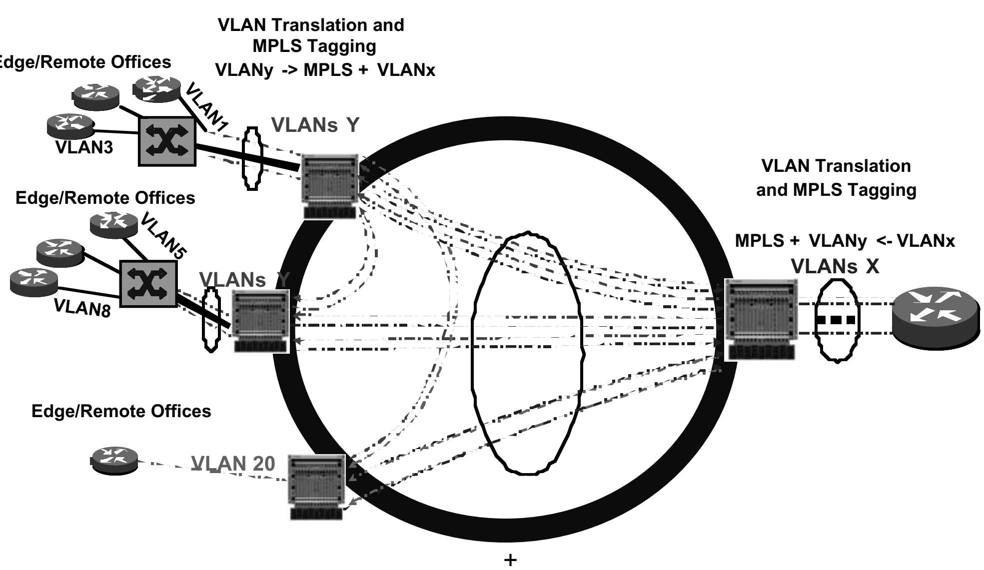

What Third-Party Developers Want

John Marshall
President and Chief Executive Officer
Intelis

Third-party applications developers write enhanced-services software for different service providers (SP). What is it that these third-party software developers seek? This paper will discuss those needs, as well as some of the products these developers might create, the role of the application programming interface (API), and why SPs would buy products from third-party developers in the first place.

Why Buy Products from Third-Party Applications Developers?

There are many third-party applications developers writing software. Why would an SP do business with such a company? SPs using a typical switch are basically playing on a level field with their competitors. They might differentiate themselves on service or price, but as far as functionality is concerned, they are all pretty much the same, meeting 80 percent of customer expectations (see *Figure 1*). Adding different services from third-party developers allows SPs, in conjunction with their switch vendors, to offer products that differentiate them from their competitors, meeting and—with the right SP—exceeding those expectations (*Figure 2*).

Third-party developers tend to focus on market niches, drilling down on a certain application or niche market, which allows them to become experts in that area. They typically have fewer customers so they know their customers intimately and can concentrate on their needs. By working so closely with the SP, they can provide products that allow the SP to differentiate itself from its competitors.

Products Available from Third-Party Developers

One application available from a third-party developer is a Web-based Centrex system that will use the Internet to control the switch. Although it still uses voice frequency over the plain old telephone service (POTS) system, the Web is used to control that user experience.

Web-based subscriber call control is another application used in conjunction with a POTS phone. When the phone rings for an incoming call, a Java applet on the Web browser flashes on the screen with the user's caller identification. The user can click a button to immediately send the call to voice mail, forward the call to an assistant or receptionist, or even let callers know that the user will be with them in a minute, allowing callers to hold or leave a message. With such applications, the Internet is used to modify the experience of the incoming caller.

Subscriber-configurable call treatment based on caller identification or time of day is another useful application available from third-party developers. Users can configure their phones to intercept calls before they ring and give the caller a choice of ringing the call through for an emergency or getting a call back at a more convenient time. For example, if a family's small children are put to bed between 7:30 pm and 9:00 pm, a call during those hours can disrupt the entire schedule. With the subscriber-configurable call treatment, any incoming calls would hear a message saying, "Hi, this is John. We are putting the children to bed right now. If your call is an emergency, press 'one' and the call will go through. Otherwise, leave a message and we will call you back after 9:00 pm." If messages have been left, the phone will automatically ring at 9:00 pm and give the subscriber the messages.

Many other services are available from third-party developers, including follow-me/find-me/single-number applications, vertical market-specific applications, and many more.

How Third-Party Developers Help SPs

By working with third-party developers, the SP can offer customers and subscribers high-margin, sticky applications that reduce churn and differentiate them from other SPs.

Third-party developers also can help SPs reduce the cost of sales by providing them with products that can be sold to targeted vertical markets. Perhaps the SP wants to have a selection of Class-5 features dedicated for real estate agents. With that type of product, competitive local-exchange carriers (CLEC) could focus their marketing and sales on real estate agents rather than on a wider audience. Having services or features for a specific vertical market considerably reduces the cost of sales.

Role of the API

The key to third-party application development is an API to interface applications to an SP's network or switch (see *Figure 3*). An API allows the third-party application to talk to the switch and essentially control or receive messages from the public switched telephone network (PSTN). The

FIGURE 1

Meeting Customer Expectations Using Functionality from a Native Switch

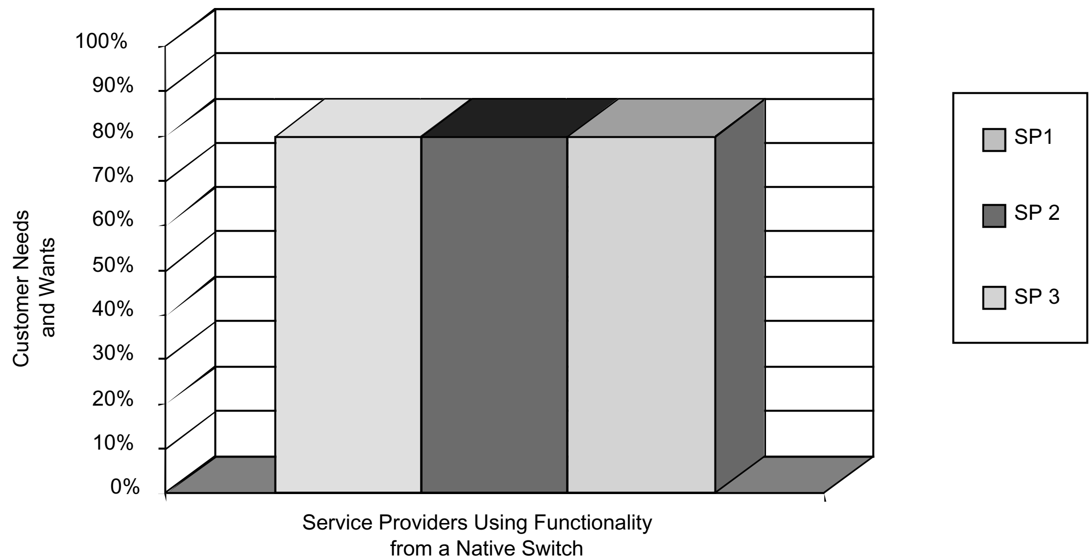

SP's customer is not interested in protocols. Customers care about the services and enhanced features available when they pick up their phones. The API is very important to third-party developers because they want to deal only with the applications they develop. Their expertise and value is in the application.

Needed API Traits

What must the API actually be able to do? First, it must be able to signal third-party applications on call events. The API can let the application know when subscribers pick up their phones or dial a number. The application may want to know what number was dialed so that information can be manipulated. Second, an API should allow the application to manipulate the dialed digits. Third, the API should allow the application to change switch parameters—change the ringing pattern or dial tone, for example. A special dial tone could indicate to the subscriber that a voice-mail message is waiting or that some special feature is activated. Different ringing patterns could be used to indicate calls from the subscriber's boss or from a known automatic number identification (ANI).

FIGURE 2

Exceeding Customer Expectations Using Products from Third-Party Developers

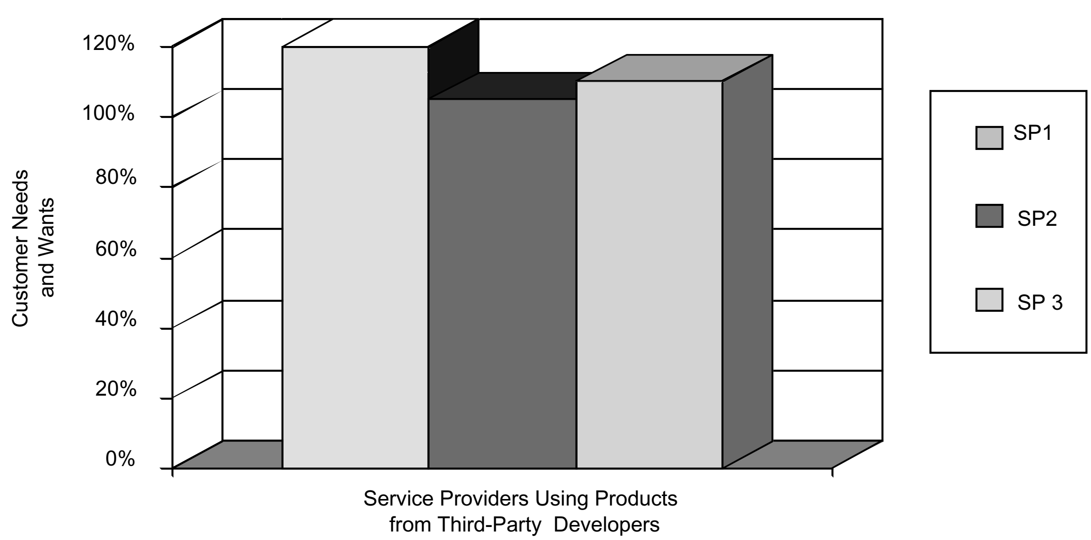

FIGURE 3

What Is an API?

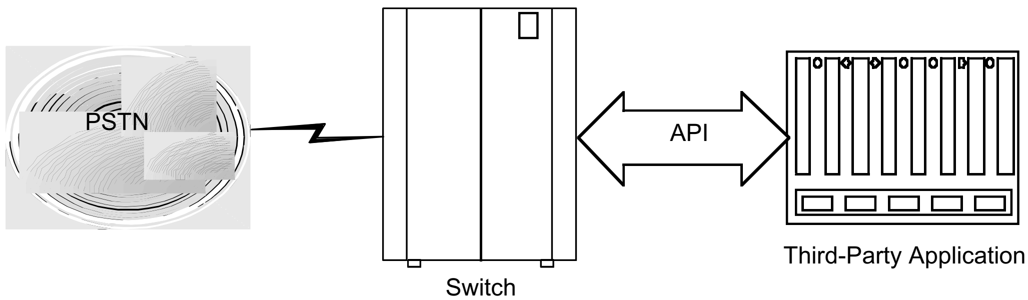

Desirable API Attributes

Third-party vendors want to see an API that works with multiple switches and multiple switch vendors. A certain amount of human effort is needed to port an application to a particular API, so the more switches and the more places the third-party developer can use to get an API, the more advantageous. The API really wants to remove the developer from the details of that call. Third-party developers are far more interested in the application as seen by the user than they are in the details of the telephony of that call. Obviously, third-party developers would love to see a rich set of features in the API. If the API does not support certain functionality, neither can the application. Therefore, the API must be robust. As new technology comes out with new protocols, that API needs to be flexible and move on with them. The API should support new technology as it becomes available, transparently if applicable.

Standards

From a third-party developer perspective, standards are great. Many standards have been discussed or are actively being developed, including session initiation protocol (SIP), Java APIs for integrated networks (JAIN), and Parlay. However, none of these standards is currently in production. Third-party developers cannot afford to wait for standards to settle down before they get into enhanced applications. How do the developers work around the lack of standards? To develop applications today, the third-party developer isolates the application from the API with a thin layer of API–dependent software. As the enhanced telephony applications are created, the various switches can be supported by modifying only this thin layer of software. When the standards settle down, the developers can simply write a new thin layer of software to interface the application to that standard.

Summary

Given an appropriate API, a third-party developer can provide SPs with applications that sell for high margins, reduce churn, and focus on vertical markets, thereby reducing the cost of sales.

For more information on this topic, please visit www.intelis-inc.com.

Building-Block Architectures Rescue Vendors Serving the Metro Optical Market

Lonnie Martin

Founder and Chief Executive Officer

White Rock Networks

Roughly $4 billion of venture capital was invested since 1995 in nearly 50 start-up companies that set out to capture their fair share of what is seen to be a burgeoning buildout of carrier metro optical transport infrastructures. Whether end-user bandwidth demand is growing 50 percent per year, 100 percent per year, or 300 percent per year, or whether there will be 20 carriers to deliver this bandwidth in the future or 200 carriers, the resultant need for optical transport equipment to eventually eradicate the metro bottleneck is very large and still growing.

New, New Networks…Or New, Old Networks?

All 50 start-ups set out to dramatically improve metro optical transport as it was practiced at the time. The largest portion of these start-ups set out to integrate the functionality of several adjacent network-element types into one system (yielding fewer boxes and fewer vendors). And the rest figured that the only way to dramatically improve the carrier's life was to invent new transport technologies altogether. The former wanted to make existing networks much better, the latter suggested that fundamentally new networks be deployed. Ever since the Internet bubble burst, and along with it the fortunes of many new carrier networks, the reality has sunk in that the financial feasibility of new networks is more important than the technical feasibility. Said another way, *evolving* existing networks to make them better has won out over *jettisoning* the old network in favor of a new network.

Yet irrespective of whether the 50 start-ups set out to deliver *new, old* networks or *new, new* networks, the architectural approach of the products deserves examination, too. And it is the purpose of this paper to discuss the traditional *integrated platform* approach and to contrast it with a new paradigm recognized by White Rock and others involving the invention of a *building-block* architecture to solve the same problems much better.

Dramatically Reducing a Carrier's Total Costs Is The Goal

Carriers must dramatically reduce their costs while simultaneously keeping up with service bandwidth growth. It is now well known that a carrier's survival will come only from dramatically lowering its cost to deliver its burgeoning data traffic and/or earning more money off that traffic. And, in fact, a principal way for start-up vendors to break into an established carrier has always been to save them a lot of money over incumbent vendor equipment—saving the carrier either capital dollars or operating dollars, or both.

How are these savings traditionally accomplished? One method often pursued is to integrate the functions of several existing systems into one system.

What metro network elements might one consider integrating? Synchronous optical network (SONET) add/drop multiplexers (ADMs), optical ADMs (OADMs), asynchronous transfer mode (ATM) switches and/or Internet protocol (IP) routers, time division multiplexing (TDM) cross-connects, and maybe service-layer processors or digital subscriber line access multiplexers (DSLAMs). And scale questions need to be answered for each of these functions. The resultant system, most often referred to as the multiservice provisioning platform (MSPP), is the equivalent of *squeezing 100 pounds into a 50-pound bag*.

The building-block architecture was invented to overcome some of the penalties of an MSPP architecture and to improve upon the result, yet do it in a *disintegrated* manner.

Metro Market Reality: It's A Diverse Application Space

From the carrier perspective, there is no single, easily described metro market to serve. The topology and customer density of metros differ as much as Manhattan and Tulsa. The end users being served differ as much as those wanting optical carrier (OC)–48 and those wanting 10BT Ethernet services, and the tough job for a carrier is never knowing where his next customer may turn up. There are fiber-rich environments and fiber-poor environments. And adding to the complexity for an equipment provider wanting to serve many carriers, the carriers differ in scale and market strategy and product offerings as much as AT&T and Covad.

In fact, the metro market space is so diverse that it is hard to imagine a single product or single approach that could cost-effectively serve most situations for any length of time. And this leads to the first penalty of a highly integrated product architecture.

Integration Penalty #1: Application Breadth = System Compromise = Sub-Optimization = Cost Penalty

The key problem of highly integrated architectures is that they compromise the cost of serving discrete transport applications to potentially serve several. That's because MSPP backplanes, control systems, and chasses must be architected to accommodate the full breadth of applications, including the so-called corner cases. This breadth has a direct impact on a system's common equipment, e.g., shelf size, backplane speeds and connectorization, power consumption, common control, etc. Common equipment has a cost, and the broader the breadth of applications, the more cost. Hence, it is an extremely rare architecture that is able to accomplish best-of-breed cost performance for any particular application over any length of time while still having the ability to address several applications.

Predictably, the eagerness and creativeness of 50 start-ups to dramatically improve metro transport for carriers by integrating all of the capability that these carriers would ever need into a single system has yielded some of the physically largest and most complicated systems ever seen in this space. Unfortunately, these systems tend to have high common equipment costs, which means that the carrier's *first cost* (the cost to deploy for his first customer) is high. Furthermore, too often the benefit of an MSPP's application breadth is of little benefit at specific sites where only one or two transport applications really exist.

In short, a systems vendor's engineering department has a very difficult system architecture equation to solve, and the trade-offs are numerous and multidimensional. Marketing wants the product to serve a very broad set of applications, and it would like the product to be the cost leader for each of them *and* for various combinations of them. To beat the competition, engineering must come up with a novel architecture, and it must decide what technologies it needs to create and what technologies it can depend upon commercially to accomplish the overall objective in a reasonable time and for a reasonable amount of money—and it must calculate the product's competitiveness two or three, or four years down the road, when the product is finally out of engineering and ready to sell. This last caveat leads to the second integration penalty.

Integration Penalty #2: Breakthrough Technology Is Hard to Predict or Count On—and That Can Kill You

Who among us can correctly predict the arrival of technologies that could make our product's functionality dramatically cheaper? On one hand, telecom systems vendors are now blessed by the so-called *horizontalization* of the telecom equipment business (as happened in the computer business in the 1980s/1990s). If venture capitalists have funded 50 optical metro transport systems companies, they have also funded hundreds of opto-electronic and optical component companies. Yet, on the other hand, this diversity of component talent and technology, and the pace of change of that technology, makes the systems architecture job very difficult, because so many decisions must be made with imperfect information. And the worst thing to do is count on a nascent technology ahead of its time.

Integration takes time—and the more that there is to be integrated, the more time it takes. If marketing and engineering try to integrate too many functions into the same systems architecture, then two or three cycles of component technology could transpire before an integrated product is finished. Most people observe that today's component obsolescence rate is nine to 15 months, and yet it may take two to three years to complete an integrated systems design. Hence, if the wrong technology choices are made at the front end during architecture design, then the resultant system could be uncompetitive by the time it's finished.

But it is potentially worse than that. Integration also adds to systems complexity—the more functions to be integrated, the higher the system's complexity. If one posits that the competitive survival of all systems vendors depends upon their ability to take regular advantage of new component technology, then the system architecture needs to accommodate this. But the more complex the system, the harder it is to go back into the system later on to dramatically improve it with new technology. Some functionality enhancements, or some cost reductions, can be done with almost all systems after they have been released, but it is very difficult to bring an integrated systems architecture back up to best-of-breed status. Rather, it usually makes more sense to start over again with a clean sheet of paper, which can be a competitive nightmare.

In general, the sheer pace of component technology advancement is forcing systems vendors into very short product life cycles and into making themselves as nimble as possible for the unpredicted technology that can offer it, or its competitors, a leapfrog advantage.

Integration Penalty #3: Application Breadth = System Size = Wasted Space (and Space Is a Precious Resource)

MSPPs cannot *do it all* in a traditional architecture without being large, even enormous, network elements. Today's breed of MSPPs fit two to four per standard 7' rack (and a couple, only 1). Space has a cost and has become a precious resource for the carrier in a deregulated world. And the worst part for a carrier is seeing a near-empty system occupying half a rack while waiting for its second or third, or fourth customer to show up. By their nature, MSPPs trade air and floorspace for application breadth. Yet this space can be as precious to carriers as dollars, and ideally, carriers shouldn't be paying for air.

Integration Penalty #4: MSPP = Complicated System = Long R&D Cycle (See Penalty #2)

MSPPs aren't easy or inexpensive to design and fully debug—even with a running start, MSPPs take at least two years just to get to Release 1.0. A lot can change in two years. Plus, it often happens that what carriers really end up wanting is the feature content of Release 2.0 or 3.0. So that means three or four years to reach a saleable product. Go back to Penalty #2. When the dust clears in three to four years, one can bank on a younger company that took advantage of the very latest component technology, producing a product with much better price/performance for those applications. And then the MSPP maker gets bound up in the following decision: Do I invest to make myself more competitive in the application, or do I cede that application to others and move on to some other opportunity (needing Release 4.0)? Eventually, the decision between reinvesting to reduce product cost for more competitiveness in today's applications, or investing to address future niches, has no other answer than starting over with a clean sheet of paper—long before the first MSPP development investment produces an ROI.

Integration Penalty #5: Long R&D Cycles + Short Market Windows = Low Probability of Decent ROI

MSPP development cycles are long because MSPPs are complicated. Market windows are short because component technologies and customer applications are changing so fast. The probability of traditional MSPPs earning decent ROIs is going down, not up. There are many reasons why start-ups and all companies can miss and/or not be competitive during a market window, but the most basic is product capability. A vendor's product strategy has to give it a chance to be nimble in order to hit, not miss, market windows. And the fewer times that a company needs to re-invent its product technology, the better the overall ROI.

Building-Block Architectures: Path to Application Breadth and Best-of-Breed?

Problem statement: How does one architect a systems product that can serve, in a best-of-breed fashion, the diversity of incumbent local-exchange carrier (ILEC), competitive local-exchange carrier (CLEC), and interexchange carrier (IXC) metro service-delivery applications in the face of component technology advances that change faster than the time it takes to finish a traditional integrated system design?

One possible answer that several companies are pursuing is to take a page from Lego, the company that popularized Lego Blocks. Can a *Lego-block family* of network elements be developed that doesn't suffer the penalties of an MSPP architecture and can provide the following?

- The same list of functions that MSPPs set out to provide
- Relatively short development cycles instead of long
- Best-of-breed manufacturing cost for most/all key metro transport applications
- The ability to significantly reduce the manufacturing cost multiple times over a product's lifetime
- A open-ended architecture with a legitimate chance to accommodate unexpected technology advances or application changes
- A means to maximize corporate ROI by stranding the least development dollars previously spent

The OSI Stack Hints at How to Break the Problem Down

The seven-layer open systems interconnection (OSI) functionality stack already hints at a building-block approach, and optical transport equipment is generally agreed to span the bottom three layers, i.e., the photonic layer, the transport layer, and the service layer as seen in *Figure 1*.

Whereas MSPPs are designed to integrate the marketing department's *optimal* mix of these three layers, a building-block family system would instead encompass multiple discrete product elements at *each* of the three layers. And these building blocks would be designed to hook together in a variety of configurations, much like Lego blocks if each was given the intelligence to provide *Lego-block glue*. In White Rock's case, this is a standards-based control plane.

Building-Block Benefit #1: Application Breadth without Platform Compromise

We can imagine several discrete building-block element examples, and we can envision just from describing them how they change the development engineering problem compared to an MSPP architecture:

- An OC–48 system doesn't need to be designed to accommodate OC–192.
- An OC–192 system doesn't need to be designed to accommodate OC–768.
- A tributary element doesn't need to be designed for both optical OC–*n* tributaries and electrical tributaries (e.g., digital signal [DS]–1, DS–3, 10/100BT).
- A coarse wavelength division multiplexing (CWDM) element doesn't need to be designed for dense wavelength division multiplexing (DWDM).
- A synchronous transport signal (STS)–1 grooming switch doesn't need to be designed to groom virtual tributaries (VTs).
- Etc., etc.

In general, we can see and appreciate that the narrower the functionality, the better able an engineering department can be to achieve a best-of-breed result for any given building-block element. And the narrower the functionality list, the fewer trade-offs and compromises that need to be made in the design of that element. The most important compromises impact the product's cost to manufacture, and if the most important goal versus competitors is to have the lowest product costs, then minimizing the number of compromises in each part of the overall product family capability is crucially important in the competitive battle.

FIGURE 1

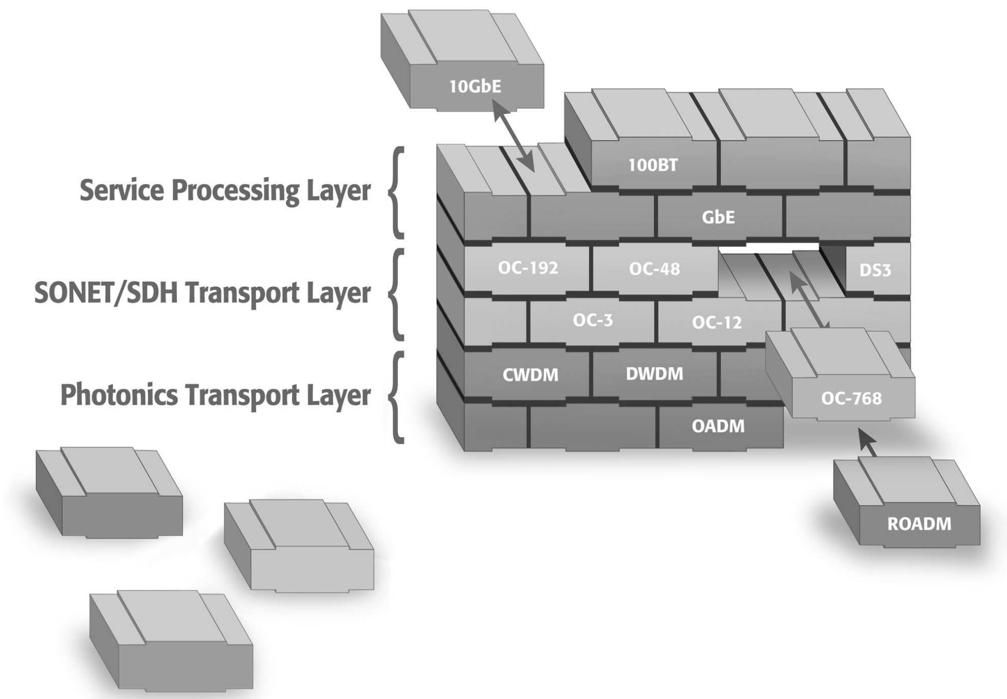

Building-Block Benefit #2: Narrower Functionality = Shorter Development Cycles = Latest Technology Benefit

While a building-block system does require several building-block elements to achieve the same overall functionality of a single MSPP, the development time required for any particular building-block element is much shorter than an MSPP. And if this development time can be reduced to nine to 15 months, then a building-block vendor has a chance to get full benefit from the latest component technologies—this is in contrast to the MSPP vendor, which is often forced into key component choices at the beginning of a three-year development.

Moreover, building-block elements represent separate developments that can be done in parallel with little conflict or interaction. By their nature, MSPP developments have numerous interdependencies between distinct parts of the system, and it is often the integration of these that is the most problematic and hardest to manage.

Provided that basic interfaces and communication links between building-block elements are architected up front, the building-block approach is showing signs of making the engineering job much easier in total.

Building-Block Benefit #3: Best-of-Breed Cost, Size, and Power

This can be a direct result of Benefit #1 if a vendor is committed to being best of breed. Anecdotally, White Rock was skeptical at first about whether a system built up from building blocks could be considerably cheaper, considerably smaller, and considerably more power efficient than the equivalent MSPP architecture. But so far the evidence supports the logic and our implementation of it, and this seems to be directly related to the compromises that must be made to develop a multifunction integrated system (see Penalty #1).

Building-Block Benefit #4: Readiness for the Unexpected Customer, Technology, or Application

Smart people try their best to predict the future, but usually fail to be perfect. The survival of a vendor in the metro optical equipment market may well depend upon its architectural and corporate ability to respond to the unexpected, and respond relatively quickly. A building-block family architecture is inherently open-ended, whereas an MSPP is too often a bounded system with hard limits that cannot be modified except by starting over. Although no vendor can react to all market events that occur, whether predicted or not, we believe that the company that is least fettered architecturally can react the fastest and can produce the best result.

The latest best example of this is the market importance of delivering Ethernet services in metro environments. One major carrier recently announced that fully 25 percent of their 2002 capital expenditure (CAPEX) budget would be spent to accomplish Ethernet service delivery. That same carrier not more than two years ago professed little interest in delivering such services. But the world has changed quickly, and now this carrier needs to change—and the vendors that supply them metro transport equipment need to react or miss out on this *sudden* opportunity. If a platform was architected two or three years ago and guessed wrong about this need, it is very likely that platform will not be able to be modified quickly, or if able to be modified, to be a best-of-breed contender.

Summary

Worldwide purchases by carriers of metro optical transport equipment will be large for a long time as the industry tries to fix the so-called metro bandwidth bottleneck. As a result, there will be heated competition from established companies and from start-up companies alike. Entering and then keeping up with the dynamism and rapid change of this market space is no simple feat for any metro equipment vendor.

The best way to serve this diverse application market from a product and technology perspective is clearly under a heated debate of sorts—otherwise, 50 new companies would not have been funded to compete with the 10 incumbent vendors that have been serving this same market for many years. Competitive advantages can be created via new technology or new product architecture, or both, and there are probably 60 opinions today on which approach may be best. Among the new architectures under development by a few of the newer companies is a building-block product family approach that shows the promise of having some important advantages over the industry's traditional integrated system approach. Time will tell whether transport equipment building blocks will catch on as well as Lego blocks have.

Meeting the Needs of Metropolitan Networks

Richard C. Notebaert
President and Chief Executive Officer
Tellabs

Abstract

The communications industry is changing, and as a communications equipment company, we must change our focus to reflect the new realities in the marketplace. Our service-provider customers have shifted their focus from investing in the newest technologies to investing in technology that manages bandwidth growth and speeds revenues. We can meet service providers' changing needs by solving the bottleneck issue in network links that are close to their customers. This article explains why metro optical networking can help service providers to lower costs, manage bandwidth, and speed revenues in metropolitan networks.

Introduction

Despite an economic downtown, end users continue to drive demand for communications services. This demand increases when companies relocate into new buildings in suburban locations, build data centers that require a large amount of bandwidth, and enhance their building with the latest technologies. Each one of these situations challenge service providers, because their customers will work with whichever company can turn up new services the quickest. And that means that service providers must devote the capital expenditures (CAPEX) needed to get high-speed services implemented and turned on faster than their competition.

Competition is strong among service providers in these situations. Some service providers go into a high-growth area early and place fiber long before companies move to these locations, giving them the advantage. In many cases, incumbent service providers need to invest new capital expenses into areas that may not produce new revenues. For example, an existing customer who relocates and needs new infrastructure may not require any more services than when at the old location. Nonetheless, the service provider is faced with making a network investment without the opportunity for increased revenue. That's why service providers need communications solutions that lower costs, manage bandwidth, and speed revenue fast enough for them to meet the needs of their customer.

Service providers are now turning their attention to ways to solve this bottleneck issue in the "last-miles" of their networks—the metropolitan network. This area of the network spans a region, such as Chicago or Memphis, or connects multiple locations in a geographical area. Service providers experience the highest and most unpredictable level of traffic in this part of the network, making it a primary focus, since it is also the area of the network with the highest growth potential. Industry experts forecast that overall traffic in the network will increase 50 percent to 60 percent annually until 2005. Aggregate voice, data, and Internet traffic in metropolitan areas will grow from \$15 billion in 2000 to \$50 billion in 2005. And, data center traffic alone is expected to grow at nearly a 100 percent compound annual growth rate (CAGR), consuming 40 percent of the total metro bandwidth by 2005.

Based on the current environment mix of increased network traffic, demand for bandwidth, and prevalent competition, service providers are challenged to meet the bandwidth needs of their customers, deliver services quickly, and with a profit in their metropolitan networks.

Delivering Revenue and Growth Opportunity

Service providers are looking for the best way to maximize their current network investment to get the most use out of their equipment. They also realize that their customers want high-speed data interconnections, which are difficult to provide without the right technology. They're confused about what is the best technology, and combination of technologies, for their company.

Technology that maximizes current network investment, integrates quickly, and delivers new services fast will win the business of these eager service providers.

At Tellabs, we believe metro optical networking will enable our customers to utilize their current network investment for traditional voice lines and upgrade to optical networking equipment to provide high-speed interconnections and services quickly. Metro optical networking will help service providers take advantage of the growth opportunity in metropolitan networks and provide valuable investment protection on legacy equipment.

Metro Optical Networking

Service providers are looking to solve the increase in bandwidth demand in metropolitan networks a couple of different ways. Some companies choose to invest in next-generation optical networking solutions, such as metro optical networking, or to focus on electronic sub-wavelength multiplexing. We believe our customers should be able to do both.

Metro optical networking enables service providers to keep signals in the optical domain, delivering new high-speed services with faster connections and increased bandwidth. It takes time to switch signals from optical to electrical, and back to optical. Metro optical networking enables service providers to send multiple signals over one lightwave using dense wavelength division multiplexing (DWDM) technology. In fact, DWDM technology allows service providers to send more than 80 different communications signals over a single optic fiber.

The lower costs that service providers realize from metro optical networking gives them room either to increase profitability or lower costs to remain competitive. In fact, next-generation technologies such as metro optical networking are 30 percent to 70 percent more cost-efficient than legacy networks.

While metro optical networking can be very beneficial, some service providers choose to focus on electrical sub-wavelength multiplexing as a way to send signals through their metropolitan networks. Sub-wavelength multiplexing keeps everything on the electronic level, enabling service providers to utilize existing network infrastructure. This saves companies money initially on network investment, but sooner than expected, the time will come when bandwidth demand exceeds available network capacity.

While most companies focus on either optical or electronic networking solutions, we believe service providers need to find the right combination of both technologies. Tellabs' metro optical networking solutions integrate easily with existing network equipment to enable companies to maximize their current network investment for traditional voice services and deliver new high-speed services quickly. These solutions also boost initial network capacity by as much as 32 times and offer the flexibility to handle today's network traffic and easily grow to manage tomorrow's optical networking needs.

Focus on Customers' Needs

Communications equipment companies can help service providers take advantage of the growth opportunities in metro optical networking by listening and respond to their needs. We need to deliver solutions that are flexible enough to allow service providers to find the right mix of electrical and optical networking that they need to deliver service and revenues quickly.

We can help service providers to speed installation, engineering, and provisioning by making fundamentally complex equipment simple. We need to develop simple engineering rules and easy design guidelines. The equipment should store the information needed to enable these rules and guidelines. An engineer that is working on the technology today needs to be able to pass the project along to new people without loosing information on how to make changes to the equipment. This effort will help service providers to integrate new technology into their existing network quicker and easier.

Conclusion

The realities of the communications industry have changed, and that means that all of us in the industry need to adjust. We need to develop solutions that address service providers' top business needs to manage bandwidth, lower costs, and speed revenues.

Service providers face their greatest challenge in metropolitan networks, where they are looking for new ways to stay competitive and generate revenues. Metro optical networking will enable service providers to integrate new high-speed optical networking equipment with existing equipment for traditional voice lines. Together, this combination of equipment enables new services to be turned up quickly and increase the capacity of existing fiber, which translates into revenue. Aside from meeting the needs service providers have today, metro optical networking also prepares networks for the future—protecting investment and lowering lifecycle costs over time.

References

1. "The Future of Metropolitan Area Networks," reported by Lehman Brothers–McKinsey & Co., August 2001.
2. "U.S. Communications Infrastructure at a Crossroads: Opportunities Amid the Gloom," reported by McKinsey & Co., August 2001.

Increasing Service Velocity in a Converged Network

Cüneyt Özveren

Chairman and Chief Executive Officer

Atoga Systems

Intelligent dense wavelength division multiplexing (DWDM) solutions are needed for the metro network so that service providers can increase service velocity and broaden their service offerings. This will enhance network performance and reduce operational costs through simplified provisioning of services. The purpose of this paper is to examine new methods for increasing service velocity and scaling in a converged network infrastructure, with specific focus on the challenges that service providers face when using this network architecture.

Network Architecture Trends

In the new market for convergence, fiber is becoming increasingly more available, and a number of drivers and applications exist that can generate profit for service providers. In 2000, the service-provider industry benefited from a competitive, capital-intensive market-share model, which increased revenue by attracting more customers. Initially, the return on investment for those companies with this type of business model was favorable. However, as bandwidth became a commodity, the market-share model resulted in an ever-extended timeline for positive return on investment and hence was no longer sustainable. Since profitability is driven by high margins, a value-based model has the best chance of driving service providers to profitability. For example, telecommunications companies have profited from the sale of relatively inexpensive local-access services. The model for cellular phones has also proved to work well for the telecommunications industry. In addition, converged network architectures are the trend, as they reduce the equipment and operations costs that arise in synchronous optical networks (SONET), DWDM, and Internet protocol (IP). Physical-network-layer convergence reduces the number of network elements, though tight integration among those layers is often difficult to achieve. Multiple protocols are still used for various purposes in converged networks, and applications have become particularly important.

Network Architecture Challenges

In order to shift to a value-based model that involves selling applications and high value-added services (VAS), companies must determine the limitations of an existing network and discover a way to migrate toward more effective network architectures. However, there are issues of time, cost, reason, and inertia to be considered in the migration to converged networks. It does not make sense to replace all existing equipment, for example. An integrated architecture can ultimately produce lower equipment costs, higher bandwidth delivery on demand, and the intelligence required to provision and deliver differentiated services (DiffServ). Some important issues include product and solution evolution, stability, standardization processes, and market acceptance. Migration also depends on the type of service provider and various economic factors. Network architecture challenges include pockets and segments with miscellaneous network architectures and overlays. Mappings and co-ordination among the layers and between networks are often quite cumbersome. Technology and protocol mappings, quality of service (QoS) mappings, and network management are potential challenges to network architectures.

End Result

The end result may lead to provisioning complications, such as a longer and more expensive provisioning cycle that requires truck rolls. Multiple network elements, increased operational costs, and complex traffic engineering complicate operations. There is also some strain on investors from the need for frequent equipment upgrades.

Service-Provider Challenges

It is often quite complicated for service providers to provide services in a converged network infrastructure. This can be accomplished, however, through protocols, tight integration between network layers, and network-resource optimization. Service requirements must be translated into the network parameters and protocol attributes. Service providers must be able to dynamically adapt to demand and network-topology changes. In the interest of QoS, it is essential to provide a wide spectrum of QoS with service-level agreements (SLA) to meet the needs of various services and applications. QoS objectives can be achieved by supporting connection-based QoS with a stringent end-to-end guarantee on loss ratio, transfer delay, and delay jitter. Connection-based QoS can be supported with a minimum of QoS requirements, allowing service providers to exploit additional network resources. Connectionless QoS meets the needs of specific services and applications through DiffServ with multiple levels of drop precedence and best-effort service. There are various access technologies in the current market, such as cable, digital subscriber line (DSL), Ethernet, and fiber. The access loop is connected to an aggregation loop that then gets connected to the backbone. The traffic must first pass through the access loop to the aggregation loop and then to a backbone, which creates too many router hops. In this architecture the multiple router hops create too much latency issues, making it impossible to provide applications to the end users.

The application service provider (ASP) model began with national service centers, though the model did not scale. With servers and services already present in regional colocation centers, service providers began to utilize local routing environments, which allowed end users to access application services directly on the optical metro network. Applications connected to an Ethernet switch within a data center can be delivered locally to the network. In this architecture, optical switching creates an incredibly fast and scalable core that pushes all intelligence to the edges. Decisions can be made at the periphery, and data can be applied to a wavelength in the optical network, assuring that the data will pass through the network quickly. The advantage lies in its exceptionally low latency. This type of architecture enables applications through high VAS, creating local environments that rely on high scalability and speed with intelligent edge devices. It is necessary to choose the right optical architecture and to build an edge device that can leverage that optical capacity to its maximum.

Provisioning

Overlay networks have a DWDM layer, a SONET layer, and an IP layer, and it is difficult to provision the various types of traffic and service offerings because there are no standards for signaling between IP, SONET, and DWDM devices. A customer who has DWDM and wants to set up an application often faces slow provisioning services that may take three to nine months, depending on whether an equipment change is needed. For example, a service provider might go out to a site and realize that there are no ports left. This situation would require the service provider to schedule another visit in order to install additional equipment that will increase capacity. This process could take six to nine months. *Figure 1* illustrates the traditional provisioning model.

Provisioning at the SONET layer often leaves responsibility to another party to make an application work within that pipe. Provisioning often works contrary to the problem that must be solved. The key issue regarding this model is that provisioning is not driven by higher-layer applications.

Application-Driven Provisioning

There are four issues to consider when driving provisioning with higher-layer applications. Overlays are separate and do not communicate with one another, which has led industry professionals to integrate those layers by collapsing them into one network element (see *Figure 2*).

Each layer is important and has a specific function. This network element combines the intelligence of the IP layer and the scaling capabilities of the DWDM layer with the reliability of the SONET layer to deliver voice, video, and data services. SONET is the best choice today for reliable transport of voice over the network without diminishing quality, and IP is a necessary component for delivering high VAS. Collapsing these layers into one box eliminates the need for carriers to manually configure three separate network layers. Information from one layer can now be used to provision the other layers. For a bandwidth-intensive application, such as videoconferencing, the ability to access that bandwidth instantly is critical. Leased optical carrier (OC)–3 is not an affordable alternative. However, since IP and SONET are in the same box, higher-layer IP intelligence can be used to trigger provisioning an OC–3 pipe instantly. This OC–3 pipe can be torn down just as quickly when the videoconference is terminated. The end customer is then charged for only the duration that the extra capacity is used. This can create significant revenue opportunity for the carrier, since such a flexible bandwidth can be charged at a premium. This is still affordable for the end user, since the capacity is paid for only when used.

FIGURE 1

Today's Provisioning Model

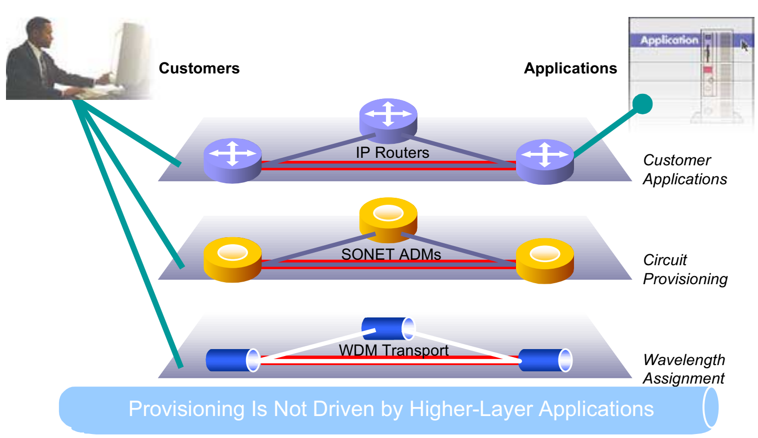

FIGURE 2

Signaling Mechanism

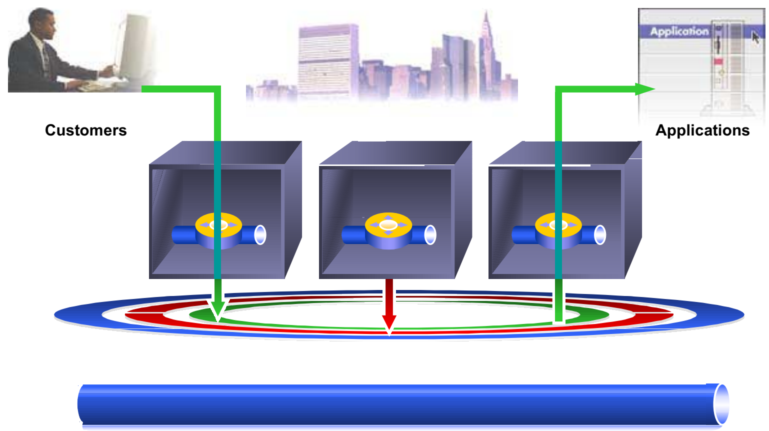

Optical Scaling

The second issue relates to the need to harness the scaling capabilities of the optical layers. SONET does not have sufficient capacity, and it is too expensive to strip out existing boxes and replace them with OC–192. Tunable lasers provide a solution to this problem, as they allow optical networks to be effectively managed and enable wavelengths to be taken from the optical backbone for data transfer. With tunable lasers, it is possible to add another card or box to the data center without performing truck rolls to other aggregation points. Instead, the lasers are retuned to create two groups (see *Figure 3*).

The network is split into two logical groups, and nodes are set to these groups, doubling the capacity. The network can be scaled to greater capacity by implementing add-in parts, and DWDM allows for greater scalability. If a request cannot be provisioned within the logical ring to which one subscribes, it is possible to look for more capacity in a different logical ring, which establishes access to the capacity available in that portion of the network (see *Figure 4*).

Figure 4 demonstrates the act of switching the node to the other logical group in order to achieve optical optimization.

Network Performance and Billing

The third issue relates to network performance and billing. Service providers who sell applications must guarantee their performance and be able to bill a customer for them. This requires SLAs on a per-user and per-application basis. Service providers must be able to monitor their applications and police the data traffic. It is possible for the architecture to assign wavelengths depending on the particular needs of applications, and performance is guaranteed precisely because data traffic is policed directly on the optical network. The act of policing traffic to a specific bandwidth

FIGURE 3

Optical Scaling

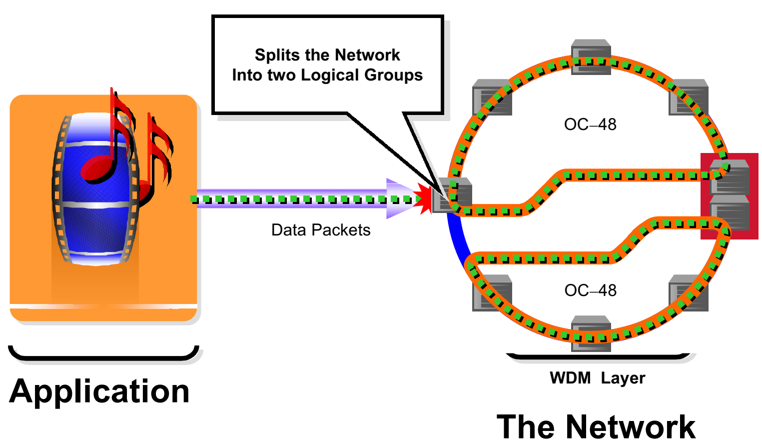

FIGURE 4

Optical Optimization

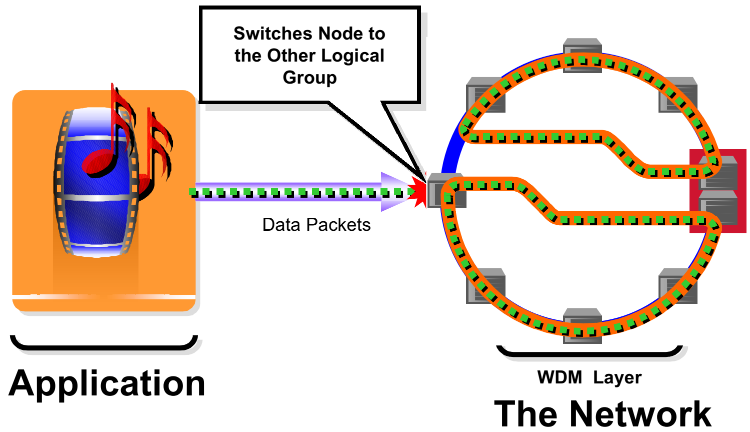

allows a service provider to provide separate billing for mission-critical applications. It is also possible to measure excess traffic that enables billing the user per usage.

Network Management

The fourth issue to be considered is network management, as it poses problems for carriers when they deal with provisioning. There are grades of service within each application, which causes database problems for carriers. Many carriers spend millions of dollars creating their own directories to solve network-management problems. Standard directories from multiple vendors are now available which perform network management. One can use a standard directory as an information repository and harness that information to manage the whole network. Directory-based product management includes good authentication features and contains applications in data sensors. Customers can access their portion of the directory in order to provision their own services instantly, add an application, or change the data service in an application. Traffic is metered and policed without delay, and there is no need to contact the service provider. This process can be activated in seconds, not months.

Protocols

Routing protocols for network reach and topology discovery are important. A Layer–2.x mechanism is employed for traffic engineering, QoS, and tunneling. Framing is possible using SONET and Ethernet, and a DWDM protocol is used for optical scaling. Converged networks offer diverse choices with regard to protection layers. The wavelength division multiplexing (WDM) layer has proprietary protection mechanisms. The SONET layer has a unidirectional path-switched ring (UPSR) or bidirectional line-switched ring. Two other protection layers include Layer 2, with a resilient packet ring (RPR), and the multiprotocol label switching (MPLS) layer. Traditional optical solutions at the metro edge usually only allow for the management of fixed transport services and typically contain Layer–1 switching intelligence. The current metro networks support applications through the use of multiple boxes and have the flexibility and efficiency to support rapidly evolving applications and services. IP has become the predominant network-layer protocol, and MPLS has allowed service providers to create IP networks that are highly scalable and capable of providing QoS.

Summary

New methods for increasing service velocity and scaling in a converged network are important, as they provide service guarantees for existing and new applications and make effective use of converged network architecture through tight integration between network layers. Service providers now have alternatives to expensive networks that require multiple devices. The combination of multiple network layers into one box reduces operational costs and simplifies network management.

Service-Provider Environment and IP Service Assurance

E. Sanders Partee
President and Chief Executive Officer
Viewgate Networks

Introduction

A primary concern of service providers today is the generation of new sources of revenue and preservation of existing revenue sources. For many of today's struggling service providers, the generation of new revenue sources often takes a back seat to preservation of existing revenue sources.

This narrow business focus results in continuation of traditional, connection-oriented data services, such as private lines, frame relay (FR), and asynchronous transfer mode (ATM). These static services are usually well managed with some level of service assurance specified within a service-level agreement (SLA). Typically, however, the assurances offered by the provider have been focused on performance of the network elements (NE) and not on the user experience of the service.

To counter this approach, corporate enterprises have implemented their own monitoring systems. Thus, they push service providers into offering service visibility and ensuring that they are actually getting the service at the contracted levels of performance. All too often, this dual approach to measuring the same service from differing perspectives produces results that do not match one another. Confusion and possibly contentious discussion erupt at quarterly provider-user account reviews, with each party defending the service performance indicated by their own set of reports.

This paper examines the area of Internet protocol (IP) service assurance from both the service-provider perspective and the customer perspective. Areas of discussion include the primary business drivers, the technical issues, actual service-delivery management (SDM) component issues, and why these topics have become increasingly important. Some business-model issues relating to service assurance are also explored.

New Service Trends

When service providers set the creation of new revenue opportunities as a top priority, they generally try to fulfill this goal by offering new services. These new services are likely to take a number of forms. For instance, there has been a shift today away from static services (connectivity-based models) characterized by monthly rentals and long-term contracts. Users are migrating towards service offerings based on episodic connectivity for use with emerging on-demand–type applications. When the nature of the application changes, so does the service. The profile of the person purchasing the service will move away from a fixed point within the enterprise, whose business is managing the telco, toward the actual business users of the service. Many of these users are scattered throughout the enterprise, not just in the information technology (IT) organization. One example could be a marketing manager who wants to do a Webcast at a particular time. Another example could be a chief financial officer (CFO), who may be renting a budgeting package over the Internetæa value-added service that would be most welcome.

These new services are changing the conceptions of what people are buying. IP virtual private networks (VPN), for example, are positioned as a replacement for connection-oriented services, such as FR and ATM. Offering additional flexibility, IP VPNs are connecting businesses and providing additional features, such as public Internet access as part of the service bundle.

The key issues then become concerns about delivery, security, and, most of all, performance. In this new environment of selling advanced services, the network is no longer the issueæthe service is. Assuring the quality of the service being purchased becomes a critical factor in the service user's buying decision and satisfaction. Assurance of these services requires measuring their quality from the user's perspective on a site-to-site basis for class of service (CoS) flow between different locations. The state of individual services need to take precedence over the state of the NEs themselves.

Technical Issues

A number of technical issues are involved in measuring these services. First, in an IP VPN scenario with the virtual router in the middle, there is a significant difference between managing the performance between different branches to each other on a single-provider network and managing performance out to another site that is not on the same network. Guaranteeing quality on off-Net participants is a difficult issue.

The second technical issue is the access network. If access to a building must go through the local phone company, managing the service quality will be limited to either the provider edge, or a reduced number of new services will be sold due to the pipe being the smallest link in the chain. From a technology point of view, the intelligent edge is both an issue and an opportunity. Currently, much of the diversity in provider edge devices, and even in customer-premises equipment (CPE) devices, has reached a manageable state. In the marketplace, however, there is an entirely new set of smarter routers and service agents that create opportunities. These new devices are able to support more differentiation for quality of service (QoS), more application awareness, and end-user–specific management systems.

These new service agents are able to deliver fine-grain QoS performance with less constraints than by traditional routers. At present, sophisticated accounting and performance measurement—for example, activating Net flow on a specific router—is difficult for carriers to do, especially on a wholesale basis. New service agents are being built to do the measurements and capture the data without slowing the router, without driving up the central processing unit (CPU) utilization, and so on. Where this measurement takes place and where these services are created—CPE versus point of presence (POP), for example—depends on what services are being built and on the size of the customer. Each has its advantages.

If the last mile to the customer is the most performance-constrained and the smallest piece of the pipe, the performance needs to be managed on both sides: both when leaving the enterprise (prioritizing that traffic and making sure the good traffic gets through first), and then from the service provider to the enterprise.

Finally, the technology is moving toward virtual private services with more on-demand services created in a dynamic environment rather than a traditional connection and security-focused VPN. Over time, IP VPNs will move to IP virtual private services, thus creating additional challenges. Traditional provisioning might be too heavyweight and unable to react quickly enough to provision a service that will last for 10 minutes or an hour and keep changing all the time. A more systematic approach may be needed to perform these on-demand services. As service-delivery speed increases, so will the requirements of service assurance.

Service Assurance Delivery

Carriers are facing several business issues in delivering service assurance. Most importantly, the customer must be kept informed. Because such services are invisible, the only way to help customers understand whether they are receiving what they purchased is to illuminate those services. A number of competitive advantages are being used in this service-assurance environment. For instance, the operations support system (OSS) can be leveraged. The typical fault, configuration, accounting, performance, and security (FCAPS) physical OSS model covers fault management, performance management, and so on. What is needed is a services OSS, which is a new breed of software.

Many larger carriers are concerned about their ability to do some of the high-touch, high–customer-focused type of tasks required with the services OSS, in which the focus is not managing devices, for example, but rather the service level and service performance. Customers are not buying the network; rather they are buying a service. And if that service begins to degrade, they will want the OSS to know about it. These customers clearly are seeking Web-based service illumination.

An environment is needed that can provide a tight loop between OSS system components, which provides active provisioning and injects low overhead. One possibility is to drive out to the customer base the ability to tune their VPNs and to tune their data services. The result is "do-it-yourself" customer service. Giving users the tools and the visibility—showing them what services they are buying and how they are performing—creates opportunities. For instance, a shared view between customer and carrier builds trust. Currently in the market of service assurance, the enterprise tool too often pushes out toward the carrier with complaints that service is not being delivered by the carrier. Bringing together a consultative tool and a collaborative problem-solving tool can build increased trust, customer loyalty, and retention. It also can provide an upselling tool for the service provider. In many cases customer unhappiness actually stems from changes in their requirements.

Customers may be trying to push too much traffic down the original circuits or the original VPN that they bought. In this case, the network is not performing poorly, and the intended service level is being met; but upgrades are needed to handle increased traffic. By discovering how well they are meeting the service needed, not just the contracted network transport, service providers may be able to generate more revenue and thus build situations that leverage the services OSS. The initial step is service design and provisioning for the first-time customers. Because a template is used—with some information from the customer about requirements—SDM and service assurance are needed to determine that delivery is meeting expectations and needs. This circles back into billing, and then the entire cycle starts again. Returning—to service the design, to redesign the system, to upsell the customer—can generate more revenue and customer retention.

Business-Model Trends

The business-model issues surrounding the selling of new services and the providing of assurance that service needs are being met are changing within corporations. The existing model with IT managers purchasing services from a carrier to support user applications is evolving towards a tendency for end users in the various departments of the enterprise to buy applications directly from providers. Part of the application that is being purchased is this services transport. As a result, the entire business model is shifting, with changes in who buys service and how quickly, what the channels are for buying it, and the marketing implications.

There also are technical driver implications. Some research on connecting facilities together has shown that the IP services market is a roughly $40 billion business. If end users could be connected to applications in a more dynamic on-demand market, however, this business would increase six fold. While the IP manager may buy a facilities connection service, today it is end users that buy the application services. The buyer is shifting from IT manager to end-user (see *Figure 1*).

FIGURE 1

End-User Purchase of Service Applications

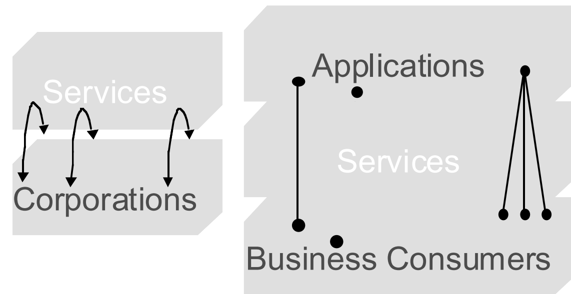

A movement from traditional implementation systems, such as traditional routers, toward more intelligent devices, such as an intelligent router called a service agent, has also been occurring. This change in the service-creation platform enables the capture of additional revenue. Consequently, the new business model has on-demand, service-appropriate bandwidth delivered in episodes. An episode in this sense might be, for example, the need for an on-demand, high-bandwidth videoconference one day a week for only 90 minutes, rather than a service purchased throughout the week. Episodes appear to be a growing trend for service delivery in the future. This trend will have relatively profound implications in the business model for service delivery.

Service retail carriers (see *Figure 2*) are emerging where organizations that are high touch can sell on-demand services, turn around and buy wholesale transport from carriers to deliver them, and be the middlemen to guarantee quality and actually put the services together. Some carriers are creating subsidiaries that can move quicker and attack this service retail market. Other carriers are creating service-creation platforms that can host these services. Still others are trying to ignore dealing with individual services.

Conclusion

The key to the new service-provider environment is to avoid becoming trapped in an outdated connectivity paradigm of simply providing facilities connection rather than actively pursuing the delivery of new services. To succeed, an OSS strategy must be developed that can provide fine-grain, customer-aware, and service-aware delivery to build and maintain leadership in the market. In addition, SDM and service assurance are critical to the new service market, which is based on charging more for delivering a value-added service.

The problem is that premium priced, value-added services cannot always be physically experienced. To overcome this, connectivity service must be illuminated so people are aware of the service's value. As the market moves on to deliver these new IP–based services, new business models and new opportunities will emerge for the forward thinking, nimble providers.

FIGURE 2

Rise of the Service Retailer

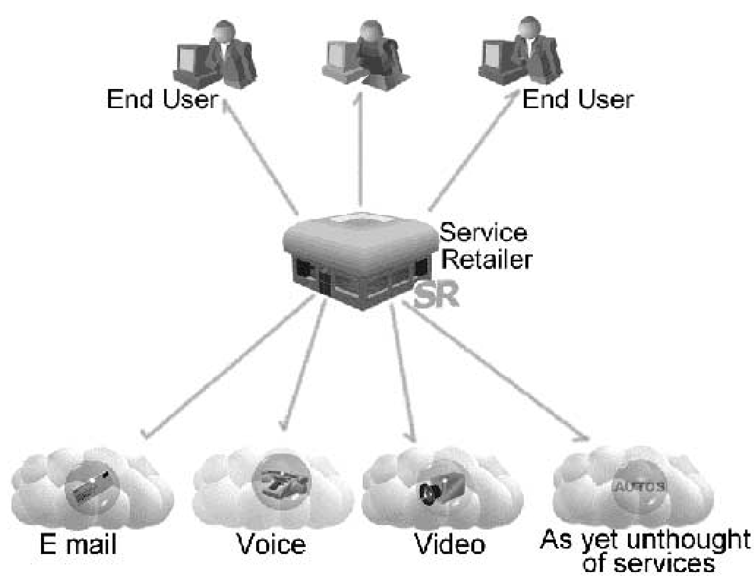

Configurable Processors Transform the System-on-Chip Platform

Chris Rowen
President and Chief Executive Officer
Tensilica Inc.

Introduction

The computer-system industry has evolved around the concept of the general-purpose computing platform: a slowly evolving instruction-set architecture, a family of hardware implementations scaling over time to higher performance levels, and a body of software packages spanning a broad range of computing applications. Starting in the late 1990s, electronics investment has shifted away from computing toward communications-centered applications, specifically to a new class of systems: high-bandwidth Internet infrastructure (data servers, integrated access controllers, core routers, Web switches) plus a vast array of focused-function end-user devices terminating the network (cell phones, personal digital assistants [PDAs], set-top boxes, digital cameras, Internet toys). The ongoing semiconductor revolution has abetted this shift and now makes it possible, and increasingly necessary, to build system-on-chip (SoC) devices that integrate all the processors, memories, peripherals, and interfaces of an entire system into a single integrated circuit (IC) platform narrowly dedicated to that application.

This article explores and illustrates the sweeping impact of this change in meaning of "platform" and highlights the potential leverage of automatic generation of application-specific hardware and software, using configurable microprocessor technology.

A Brief History of Platforms

The concept of a computing platform has powerful technical and economic motivations. Broad use of a common programming model, hardware implementation, and an applications portfolio has the potential to amortize the cost of initial development across a great volume of products and to reduce the cost and risk of new development. IBM's System 360, introduced in 1964, is considered the first major platform architecture, but each major generation of electronic systems has given birth to a dominant platform—for example, the Digital Equipment VAX introduced in 1976 and dominant into the 1980s, and the IBM/Intel/Microsoft PC, introduced in 1981 and dominant up to the present. Each of these platforms supported programming models supporting a wide range of applications (from payroll to mechanical design), and the core hardware implementation consisted of multiple chips (processors, memories, peripheral controllers), many of which were standard components intended for broad use outside the platform. In most cases, these platforms were designed or evolved over a period of years, accumulating hardware features and software infrastructure to support the broad set of target applications. Moreover, each platform gained flexibility by incorporating a general-purpose microprocessor architecture supporting a broad set of generic data-types—especially integers, character, Boolean and floating-point data.

The Design Paradox

Two fundamental shifts—one technical and one economic—are changing the design of platforms. First, continuing growth in silicon-chip capability—Moore's Law—is rapidly reducing the number of chips in the typical system and increasing benefits in cost, bandwidth, and power for the integration of multiple functions together on a single chip. Systems built from large collections of generic ICs, combined at the printed circuit-board level, are less and less able to compete technically with highly integrated systems that combine their processors, memories, peripherals, and other logic together on one or a small number of ICs. Second, growth in the value and importance of general-purpose computer systems (PCs, workstations, mainframes) has slowed significantly, in favor of a explosion of thousands of new types of electronic platforms aimed at communications and entertainment—network routers, MP3 players, cell phones, PDAs, and other focused functionality devices. Moreover, the global market for these devices is intensely competitive, so low manufacturing cost is key. These are highly flexible, yet very lean, new platforms.

This creates the central paradox of modern platform design: How do architects develop all of these new platforms that combine the key benefits—longevity, amortization of development costs over large volume, adaptability to changing market requirements—without taking too long or creating something so general-purpose that it fails to meet the brutal cost requirements of this new type of electronic platform. Without a sharp reduction in the cost of development of these new platforms, the variety and efficiency of new platforms will be significantly constrained.

How to Build a Platform

We could resolve the paradox if we could take all of the phases of normal platform design and compress them:

1. Compress the time and effort to develop, implement, and verify new microprocessors, including appropriate new instruction-set architectures, compilers, operating systems, and libraries optimized for the class of applications that will run on the new platform.
2. Compress the time and effort to develop, implement, and verify new system architectures that incorporate these processors and their surrounding interfaces and programming models.
3. Compress the physical implementation of the system—processors, memories, IO, and other logic into dense integrated circuits.

Compressed creation of new microprocessors for these lean platforms is, perhaps, the most daunting of these tasks. Traditionally, the time-scale for the development of substantially new processor architectures is years, especially when the development of new software development environments (compilers, debuggers, verification tools), and the porting of operating systems are included. Reducing the time and cost of development of a new processor accelerates the whole platform development and enables a much wider community of engineers to be platform architects. Configurable processor technology is the key.

Configurable Processors

Configurable processor technology focuses on the automatic generation of new platform-specific processor hardware design and software tools from a single high-level description. By automating hardware and software generation, the platform designer is guaranteed consistency among all of the representations of the processor definition—the hardware design, test benches, simulation models, compilers, assemblers, debuggers, and real-time operating systems (RTOSs) are all built for exactly the same architecture. This potentially eliminates the "Tower of Babel" that surrounds many embedded processor environments today, where subtle differences among members of a processor family—differences in instruction-set variations, memory system organization, debug facilities, and processor control structures—frustrate the system designer putting together the hardware, verification, and software tools. By generating the processor from a high-level description, the platform designer gets control over all of the relevant cost, performance, and function attributes of the processor subsystem, without having to become a microprocessor design expert. This effectively opens up processor design to a broad population of system and application architects, just as the proliferation of application-specific integrated circuit (ASIC) and logic synthesis tools democratized IC design during the past decade.

The four key questions for configurable processors are these:

1. What target characteristics of the processor can be configured?
2. How are the target characteristics captured by the platform designer?
3. What are the deliverables—the hardware and software components—to the platform designer?
4. What are typical results for building new platforms to address emerging "post–PC" communications and consumer platforms?

What's Configurable?

The notion of adapting the processor to fit the application is increasingly popular. A wealth of terms has appeared to describe the rich set of possibilities. The phrases "configurable processor," "extensible processor," "reconfigurable processor," and similar terms have only entered the electronics lexicon during the past three or four years, so it is useful to roughly define some of the terms.

Synthesizable Processor: A wide variety of processor architectures have been implemented in RTL to improve process portability and ease integration into standard ASIC design flows that include the generation of gate-level net lists from RTL by logic synthesis tools such as Synopsys's Design Compiler. However, synthesizability does not imply configurability. Moreover, architectures originally tuned for custom circuit implementation may require significant redesign to fit the stricter clocking, power management, and data-path design constraints of synthesis. Many popular architectures—including asynchronous response mode (ARM), millions of instructions per second (MIPS), Xtensa, and Motorola ColdFire—have been implemented as synthesizable cores.

Configurable Processor: A configurable processor is a design and tool environment that enables significant adaptation by changing major processor functions to tune the processor to specific application requirements. Typical forms of configurability include additions, deletions, and modifications to memories, to external interface widths and protocols, to commonly used processor peripherals, and, most importantly, to instruction-set functionality and execution pipelines. To be useful, configuration of the processor must meet two important criteria:

1. The configuration mechanism must accelerate and simplify the creation of legal and useful configurations.
2. The generated processor must include complete hardware descriptions (typically RTL descriptions such as Verilog or VHDL), software development tools (compilers, debuggers, assemblers, and RTOSs), and verification aids (simulation models, diagnostics, and test support).

Note that the concept of a configurable processor architecture is different from the idea of a reconfigurable processor IC described in a following section. Configurable processors may be implemented in many different hardware forms, ranging from ASICs with hardware implementation cycles of many weeks to field programmable gate arrays (FPGAs) with implementation cycles of just minutes. An important subset of configurable processors is extensible processors—processors whose functions, especially instruction-set, can be extended by the application developer to include features never considered by the original processor designer. Automated software support for extended features is especially challenging, however.

Application-Specific Processor: This is the broadest and most generic term for processors tuned for various embedded application classes. This may include both processors

FIGURE 1

Instruction Set	Memory System	Interface	Processor Peripherals
• Extensions to ALU functions on general registers (e.g., population count instruction) • Coprocessors supporting application-specific data types (e.g., network packets, pixel blocks) • High-performance arithmetic and DSP (e.g., compound DSP instructions, vector/SIMD, floating point) • Selection among function-unit implementations (e.g., small iterative multiplier versus pipelined array multiplier)	• Instruction cache size, associativity , and line size • Data cache size, associativity , line size, and write policy • Memory protection and translation (by segment, by page) • Instruction and data RAM/ROM size and address range	• External bus interface width, protocol, and address decoding • Direct connection of system control registers to internal register and data ports • Mapping of special purpose memories (queues, multiported memories) • State visibility trace ports and JTAG–based debug ports	• Timers • Interrupt controller: number, priority, type, fast switching registers • Exception vectors addresses • Remote debug and breakpoint controls

developed by traditional manual methods and automatically generated configurable processors. Examples of manually developed application-specific processors include the ARM9E (ARM with basic digital signal processor [DSP] extensions), MIPS's application-specific extension (ASE) for graphics, and Equator's MAP-CA general-purpose media. These designs are typified by multi-year development cycles and manual adaptation of software tools. The greater development cost and time for manual adaptation typically mean that these designs cannot afford to be narrowly application-specific and may not be able to achieve high application efficiency.

FPGA–Based Processor: As the capacity of FPGAs has grown, they are gradually changing their role from glue-logic and special-function block to full system implementation vehicle. Almost by definition, this means that processors are starting to appear within devices that contain lots of field programmable logic. Three categories have emerged:

- *Integration of Core Plus FPGA:* A standard fixed architecture core is incorporated on the same die as field programmable logic. Altera's Excalibur family and Triscend's products fall into this category. The processors are not field adaptable and the field programmable gates are used only for peripheral and interface functions.
- *Reconfigurable Processor:* An important subset of the reconfigurable processors synthesize parallel DSP architectures into specialized FPGA blocks controlled by reduced instruction-set computing (RISC) cores to create specialized field-configurable platforms often dedicated to complex communications infrastructure applications. Examples include the Chameleon CS2000 Reconfigurable Communications Processor (RCP) and MorphICs. Similar reconfigurable processors use small programmable engines or special DSP datapaths for executing parallel DSP algorithms, often in conjunction with RISC control cores. Examples include Broadcom's Calisto family (based on Xtensa) and the UC Berkeley Pleiades low-power heterogeneous reconfigurable multimedia processor project.
- *General Configurable Processors Implemented in FPGA:* Automatic generation of efficient processors of configurable architectures complements high-capacity FPGAs. In these cases, high-end FPGAs, such as Altera APEX and Xilinx Virtex devices, are used as synthesis targets. Tensilica's Xtensa and ARC's Tangent are both commonly used in this form. A variety of 8b and 16b processors are also available with some configurability features, including Atmel's AVR 8b RISC, Altera's NIOS, and several 8051 variants. There are no major differences between processors suitable for FPGA versus ASIC implementation—all designs want processors that are small, fast, and fully supported by tools. In practice, the processors in FPGA are among the smallest cores available and are used in applications that demand fast time to market but can tolerate the clock-frequency and unit-cost penalties of FPGA implementation.

Configurability of Processor Features

The goal for configuability is to allow features to be added or adapted in any form that optimizes the cost, power, and application performance of the processor. In practice, this can be broken into four categories (see *Figure 1*).

A color-coded block diagram for a configurable processor is shown in *Figure 2*, identifying baseline instruction-set architecture features, scaleable memories and interfaces, optional and configurable processor peripherals, selectable DSP coprocessors and facilities to integration user-defined instruction-set extensions.

The basic flow for generating a processor is shown in *Figure 3*.

The chip designer, the application expert, comes to the Web-based generator interface and selects or describes the instruction-set options, memory hierarchy, closely-coupled peripherals, and external interfaces required by the application. The generator produces both the complete synthesizable hardware design and the software environment in less than an hour. The hardware can be immediately integrated

FIGURE 2

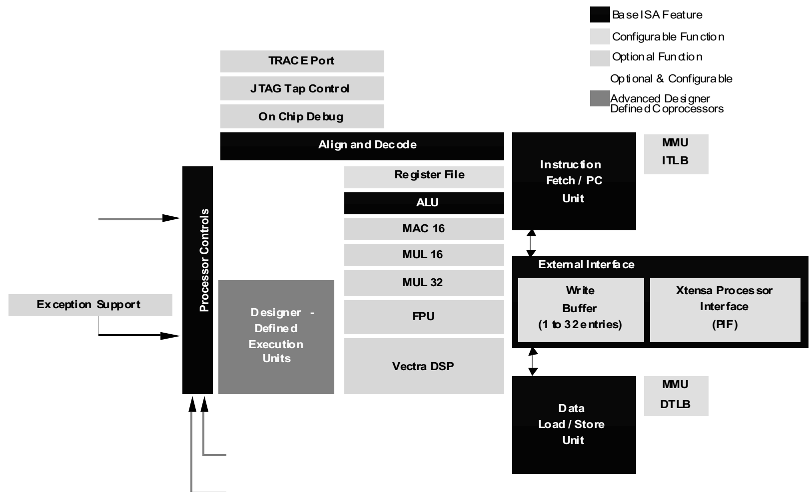

into the balance of the SoC ASIC design. It is easily and legally portable to any process or fab, ensuring optimal performance and silicon leverage. The software development and tuning can also start immediately. With the integrated profiler and one-hour turn-around for software tools and RTOSs, the designer can, for the first time, realistically tune the processor to fit the application.

How Is a Processor Described?

Fast and easy configuration depends on an appropriate means to describe and modify processor functions, especially instruction-set functions. Academic researchers have created instruction-set description languages, especially for the teaching, documentation, and generation of software tools and performance models, plus some limited automatic generation of hardware. The "Tensilica Instruction Extension" (TIE) language is the first widely used commercial instruction extension format and processor extension generator.

The TIE Language

A TIE description consists three basic parts:

- *State Declarations and Types:* Designers may add state registers and register files of any width and number. New C data-types can be associated with new register files.
- *Instruction Encodings and Formats:* Designers may specify new formats with up to six source and destination registers, including encoded immediate fields. Each new instruction gets a unique encoding.
- *Instruction Semantics:* For each group of instructions and data-type, designers specify the corresponding transformations from source registers to destinations registers. The designer may optionally specify pipeline latency, used by the compiler, simulator, and hardware

FIGURE 3

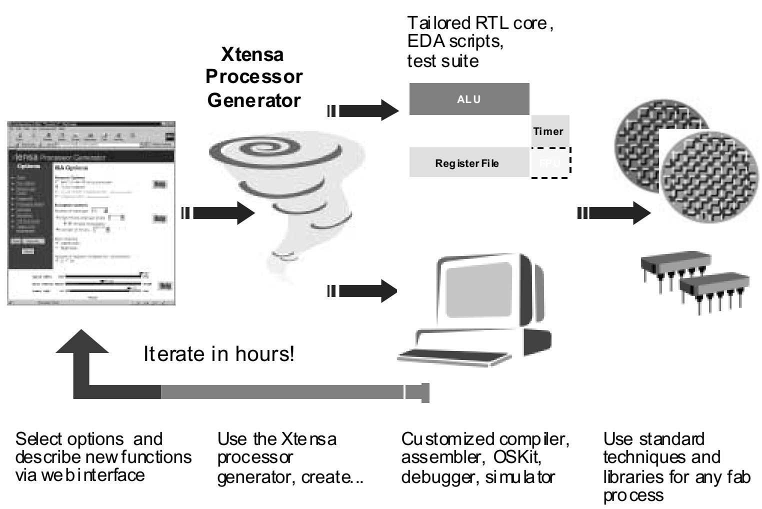

generator to automatically stretch complex functions across multiple clock cycles.

A Simple TIE Example

Here is a simple, but complete, TIE example of adding a set of instructions to the processor. It defines a new C data-type, "long128," and associates with it a 16-entry register file, where is each entry is 128b wide. The instructions that use the new register file include loads and stores, by default, plus one arithmetic instruction "add128." This instruction is pipelined so the 128b operation does not impact the clock frequency of the processor.

```
;use generator;
;generator::regfile   (name =>  "L",
;                      cname =>  "long128",
;                      sname =>  "s",
;                      width =>  128);
opcode add128 op2=0 CUST1
iclass ipv {add128} {out sr, in ss, in st}
reference add128 { assign sr = st + ss;}
schedule c2 {add128} { use ss 1; use st 1; def
sr 2;}
```

The software developer might directly use the new data-type and operations as follows:

```
main() {
int i;
long128[256] source1,source2,destination;
   for (i=0;i<256;i++} destination[i] =
add128(source1[i],source2[i]);
}
```

What Are the Deliverables?

TIE–based extensions are created by application experts, typically starting from the application kernel. They use Tensilica's baseline simulation tools to profile the application and identify hot spots in the original code, and then describe new data-types and operations that are more natural to the problem. All hardware, verification infrastructure, and software tools are created together from a single configuration file and delivered securely over the Web so that the designer can typically evolve an optimized configuration over a few days. A typical flow is shown schematically in *Figure 4*.

From that single configuration description, the processor generator produces the fast cycle-accurate instruction-set simulator and bus-functional model; complete C source libraries for the extensions; the full GNU–derived C/C++ compilers, including modified front-end, code-generator, code-scheduler, and intrinsic support; the OSKit, which is the support for kernel and debug extensions for third-party RTOSs; the hardware test-bench and complete diagnostics, including coverage of TIE extensions; the scripts for synthesis, place and route, test generation, and timing and power optimization and simulation; and, of course, complete source code in Verilog and VHDL.

What Is the Impact?

The promised benefits of automated processor generation are significant improvements in application efficiency (throughput, power, silicon cost) combined with simplification of full platform development (completeness and immediacy of compilers, simulators, RTOSs and hardware). To illustrate, it is useful to walk briefly through examples in multimedia, telecommunications, and network processing.

Consumer Multimedia

Video processing lies at the heart of consumer electronics—in digital video cameras, digital television, and games. Common tasks include color-space conversion, two-dimensional filters, and image compression. The industry-standard EEMBC "Consumermark" benchmarks

FIGURE 4

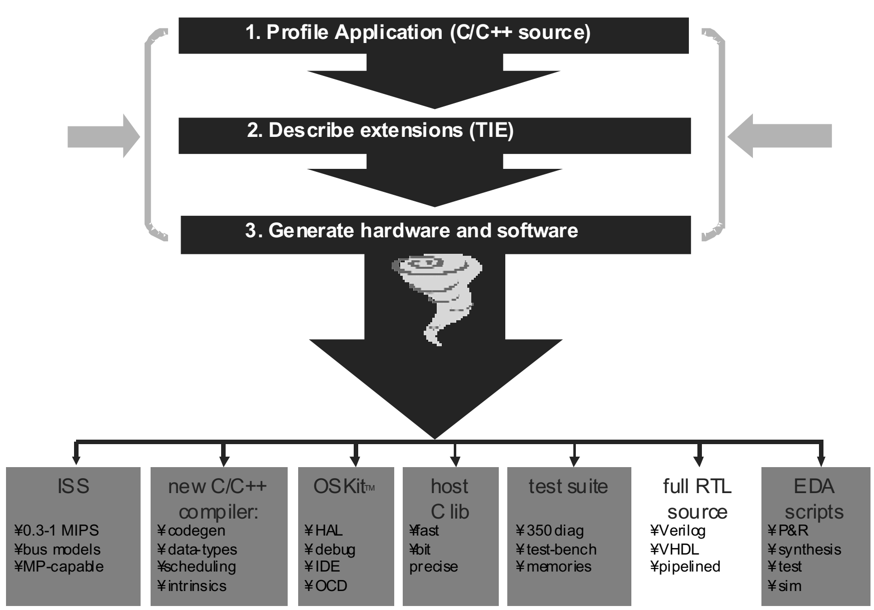

FIGURE 5

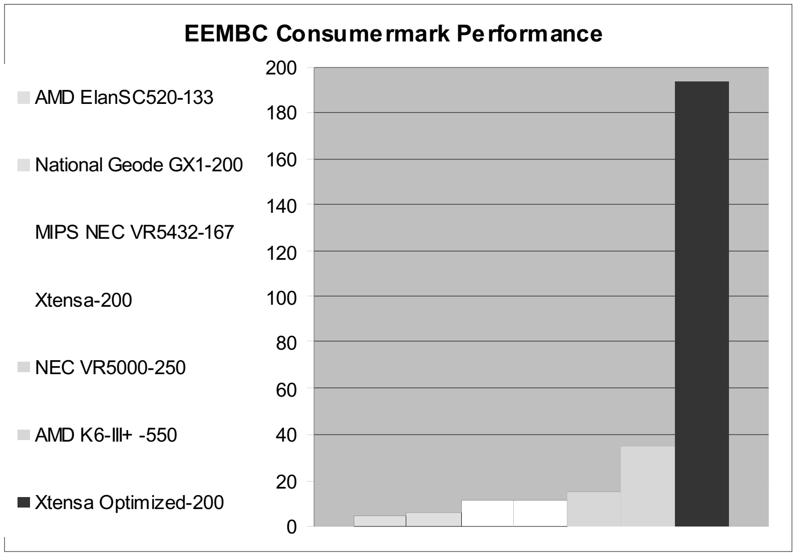

include a representative sample of all these applications. A standard configuration of Tensilica's Xtensa processor already includes many appropriate features for these tasks, and even this baseline configuration at 200 MHz (typical for 0.1µ complementary metal oxide semiconductor [CMOS] technology) delivers performance more than 11 times the performance reference, the ST20C2 at 50 MHz. However, when instructions for image filtering and color-space conversion are added using TIE, the average performance is increase by a further 17 times, resulting in a processor with almost 200 times the performance of the reference processor as shown in *Figure 5*.

These performance levels have previously only been achieved by specialized non-programmable logic blocks dedicated to a specific imaging function. The TIE add-ons include special instructions for gray-scale filtering and RGB-YIQ and RGB-CYMB color-space conversion.

FIGURE 6

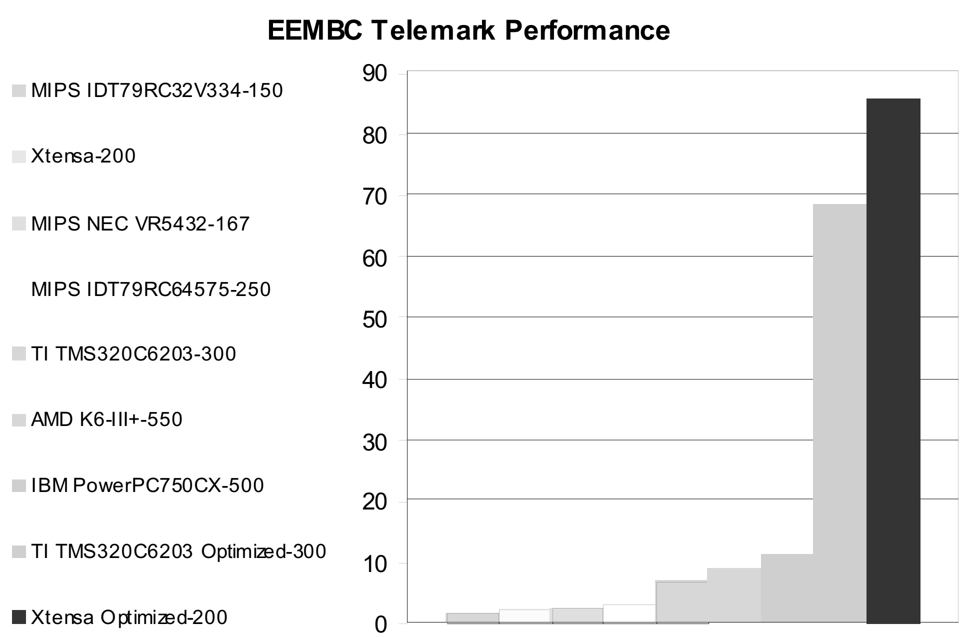

DSP Acceleration in Telecommunications

Telecommunications applications present a different set of challenges. Here, the data is often represented as 16b fixed-point values, as densely compacted bit-streams, or as redundantly encoded data. During the past 10 years, standard DSPs have evolved to address many filtering, error-correction, and transform algorithms. The EEMBC "Telemark" benchmark includes many of these common tasks. In this case, the applications designer might start with an Xtensa with the Vectra DSP add-on—a family of vector (single instruction multiple data) coprocessors tuned to digital signal processing. This gives baseline performance on the EEMBC benchmarks that compares well with other leading 32b and 64b RISC processors. However, when some additional instructions are added using TIE and the code is re-optimized to exploit the configuration, the performance jumps another 37 times. The overall performance then exceeds that of the previous performance leader, a high-end Texas Instruments C62x VLIW DSP using hand-optimized code, as shown in *Figure 6*.

Conclusion

Innovation in embedded systems is almost frightening in its pace and diversity, particularly as product creators push to deliver both more flexibility and ease of use, as well as fundamental breakthroughs in bandwidth, battery life, and cost. On the surface, it would appear than ever-greater software functionality and flexibility stands in conflict to ever-greater hardware specialization. The conflict can be resolved by taking a fresh look at how these new SoC platforms are defined, designed, verified, and programmed so that a larger number of more application-specific platforms can be safely created by a wider population of engineers. Automated generation of application-specific platforms around configurable processors is one key to this impending explosion of new platforms.

Superconductor Systems for Wireless Networks

Randy Simon
Vice President and Chief Technical Officer
Conductus, Inc.

After more than a decade of materials and applications development of high-temperature superconductors for electronics, wireless communications has emerged as the market with the greatest potential. The explosive growth of the wireless industry coupled with the increasing technology demands of advanced wireless networks have created a significant opportunity for high-temperature superconducting technology.

The focus for using high-temperature superconductor (HTS) technology in wireless applications is in the form of front-end systems for base stations that exploit the properties of superconductors to create high-performance radio frequency (RF) filters coupled with the properties of cryocooled semiconductors to enhance pre-amplifier performance. The resultant receiver front ends simultaneously offer enhanced sensitivity to improve signal reception and exceptional selectivity to reject interfering signals. The underlying superconductive technology supporting wireless applications has been well enough in hand for several years to permit the intense product development required for widespread deployment. There have already been several generations of superconductor wireless systems that include a growing number of features and performance benefits. The evolving demands of high-performance digital wireless systems continue to drive efforts to further advance superconductor filter technology and enhance its ability to address the future needs of the wireless industry.

This technology has a unique role to play in wireless cell sites because the properties of superconductors make it possible to create filters that simultaneously deliver optimum performance with respect to adjacent band rejection and insertion loss. This is because superconductors are nearly lossless at microwave frequencies. As a result, extremely sharp filters—ones with a large number of poles or filter stages—can be built without incurring substantial insertion losses, whereas filters made from ordinary materials become increasingly lossy as they become more selective. Filters made using superconductors provide the closest approximation to a perfect filter—namely, one that allows 100% of the desired signals to pass through and rejects 100% of the unwanted signals. Hence, superconducting filters are ideally suited for rejecting out-of-band signals, particularly those that are very close in frequency to the desired band. This performance is exactly what is needed for preselect filters in wireless base-station receivers.

The ability to reject adjacent band interference is important in the base station because such signals can otherwise saturate the amplifier or introduce other nonlinearities in the receiver front end and introduce distortion into voice channels, reduce capacity in digital systems, and degrade data transmission. Superconducting filters can allow service providers to utilize their available spectrum without introducing guard bands or blocked channels at the band edge. By using extremely sharp adjacent-band rejection filters—commonly referred to as "brick-wall" filters—it is possible to maintain a greater number of channels in the band. In digital wireless systems, high selectivity against interference reduces the effective noise floor, thereby restoring lost capacity and reduced coverage as well as permitting lower mobile transmit power (see *Figure 1*).

The other major benefit of having superconductor and, more generally, cryoelectronic technology in base stations is the ability to increase the sensitivity of the receiver, thereby providing extended range and enhanced coverage. The sensitivity enhancement comes about by virtue of reduced noise in the receiver. The low noise results from two different effects: the nearly perfect efficiency of the superconducting filters plus the use of specially designed low-noise amplifiers (LNAs) whose own noise levels are reduced by operating at the same very low temperatures required for the superconducting filters. Like many other electronic components, properly designed semiconductor LNAs deliver enhanced noise performance when operated at reduced temperatures. Co-locating these LNAs with superconducting filters provides an ultra-low–noise receiver front end. Having a very low noise figure increases range and coverage, reduces mobile transmit power, increases capacity, and increases data throughput for 2.5G and 3G systems. By eliminating the conventional trade-off between sensitivity and selectivity, superconductor front ends give the service provider the ability to maximize coverage, capacity, quality, and data rate all at once.

FIGURE 1

By deploying a ClearSite superconductor system at each base station, cell-site coverage increases, eliminating the dead zone without adding a new cell site.

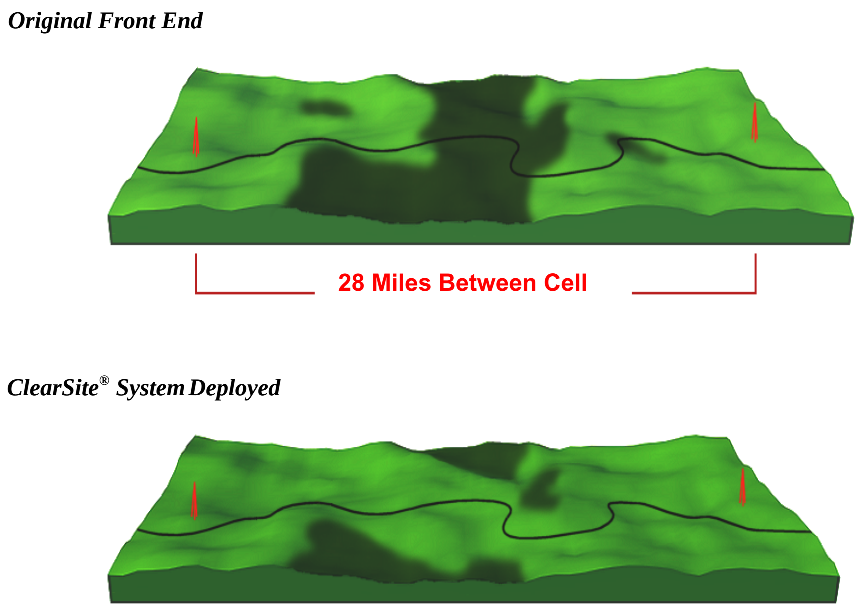

Source: Actual Carrier Field Data

The first field tests of superconductor front ends for wireless networks took place in the 1995–1996 time frame. For the first few years, when analog networks were dominant, the primary application for the technology was coverage expansion in rural networks primarily associated with the enhanced sensitivity of the front end. Over time, the application focus has shifted toward the more demanding requirements of digital networks in urban and suburban areas, and the opportunities for widespread deployment have expanded. Currently, three companies in the United States—Conductus, ISCO International, and Superconductor Technologies—are selling superconductor front-end systems to the wireless industry. As of mid-2001, several of the top-ten U.S. carriers have begun to deploy these front-end systems in their networks, and there are well over 1,000 of them in use across the country. These deployments have clearly demonstrated the ability of these products to increase minutes of use (MOU), reduce dropped call rates, and improve quality of service (QoS). For example, Dobson Communications has deployed ClearSite® systems from Conductus, and they have seen an increase in MOU by up to 30 percent in cell sites where the systems are installed.

More recently, extensive second-generation (2G) and third-generation (3G) urban field trials have clearly demonstrated the ability of superconductor systems to reduce the effects of interference on high-performance networks. In particular, the systems were able to restore capacity and coverage in cell sites that were dramatically degraded by the effects of interference. In addition, the systems significantly increased data rates and reduced bit error rates. From all indications, the benefits of these systems increase in higher-performance wireless networks (see *Figure 2*).

Urban cell sites operate in an increasingly hostile RF environment due to the proliferation of wireless service, the increases in wireless traffic, and the growing use of co-location. One particular source of interference that is increasingly prevalent in the United States is Enhanced Specialized Mobile Radio (ESMR) sites, as much as 50 percent of which are located within a kilometer of urban cellular sites. These sites produce strong interference signals very close to the cellular frequency bands.

While the effects of interference are increasingly evident in modern digital wireless networks, identifying and cataloguing the sources of interference is not a simple matter. In call division multiple access (CDMA) systems, for example, it is difficult for the switch to separate interferers from the desired CDMA carriers. The interference sources are typically variable and transient in nature. As a result, complex measurements and data reduction are required in order to accurately gauge the nature and impact of interference and to tie it to key performance metrics, such as system capacity and mobile transmit power. But despite these technical issues related to quantitatively understanding the nature of interference in wireless systems, there is a growing body of data demonstrating the impact of interference on digital networks.

CDMA systems are particularly vulnerable to interference by virtue of the unique power-control aspects of the technology. The power-control system causes all mobiles to increase power when interference signals are present in the band. These interferers can be in-band signals from numerous sources, or they can be intermodulation products caused by out-of-band signals—typically from either co-located or nearly co-located transmitters, or from competi-

FIGURE 2

The ClearSite superconductor system reduces interference at an urban TDMA cell site, resulting in an 80 percent capacity improvement.

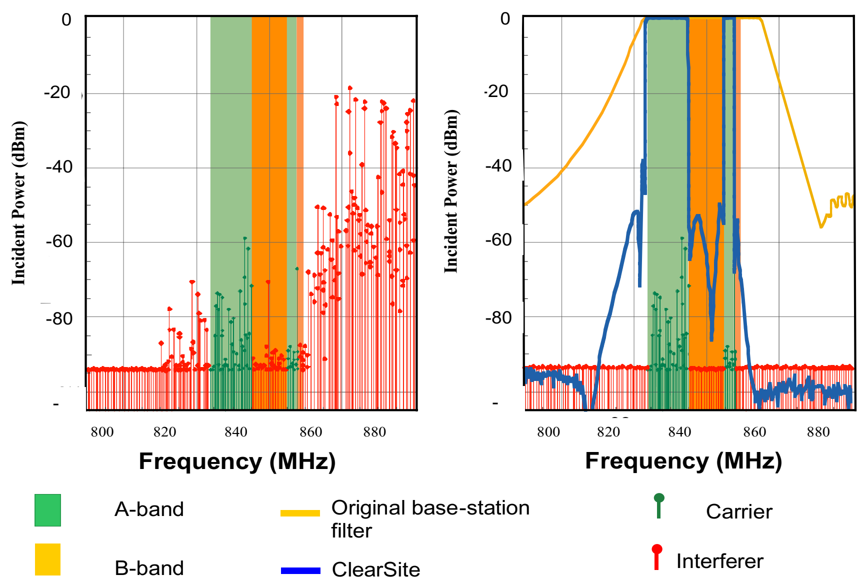

Spectrum analyzer measurements at urban cell

tive mobiles far from their cell site and thus operating at maximum power. The effect of this power-control activity can go beyond the usual "cell breathing," which occurs as more and more subscribers use a carrier, and dramatically reduce coverage and capacity.

These effects and others have been studied extensively in field trials of superconductor systems. Conductus performed one such trial with a suburban CDMA B-band operator that was experiencing interference problems from a competing A-band operator. The goal of the trial was to quantify the capacity improvement that the HTS system could provide by diminishing the effects of the interference. Calibrated interference signals were applied at levels comparable to the measured interference at the site. These signals were observed to essentially collapse the coverage area of the site using only the existing front end. With the superconductor system in place, the full coverage and capacity of the cell site were restored (see *Figure 3*). Trials in time division multiple access (TDMA) systems have demonstrated results of comparable significance.

As wireless networks move toward increased utilization for data communications and 2.5G and 3G technologies are employed, the benefits of superconductor front-end technology will increase. The ability of wireless systems to support high data rates is directly related to the signal strength to interference ratio, as described by Shannon's Theorem. By reducing noise and diminishing interference, superconductor front ends are very effective in increasing the maximum data rate in wireless systems.

Field trials of 3G systems using the International Mobile Telecommunications (IMT)–2000 wideband CDMA (WCDMA) technology have demonstrated an even greater vulnerability to interference than was seen in standard CDMA systems. Conductus performed a field trial with KDDI and Hitachi in Japan that demonstrated the effects of interference caused by the Personal Handiphone System (PHS) and the ability of the superconductor system to counteract these effects. In this trial, the 3G cell site was again observed to essentially collapse in the presence of frequent but moderate-level interferers (See Figure 3). The superconductor front end fully restored capacity, coverage, and data rate, and improved bit error rate by three orders of magnitude.

The growing installed base of superconductor front-end systems in the wireless infrastructure as well as impressive results from a variety of field trials being conducted by all vendors provide an increasingly persuasive argument for widespread deployment of the technology. Alternative approaches to base-station enhancement, such as tower-mounted amplifiers and smart antennas, have distinct disadvantages when compared with the superconductor solution, including increased vulnerability to interference and cost, respectively. There are a number of criteria that superconductor systems must meet in order to gain widespread acceptance throughout the industry. From a performance perspective, of course, the products must provide a combination of unparalleled high selectivity as well as low noise (high sensitivity). These characteristics are the reason for using the technology in the first place. But beyond the electrical performance, the systems should be highly reliable, compact, and field-proven, as they are replacing existing technology that is very reliable and has generally not been a focus of much attention in the base station. Because the superconductor systems include a mechanical refrigeration

FIGURE 3

The ClearSite superconducting system reduces out-of-band interferers at a 3G site.

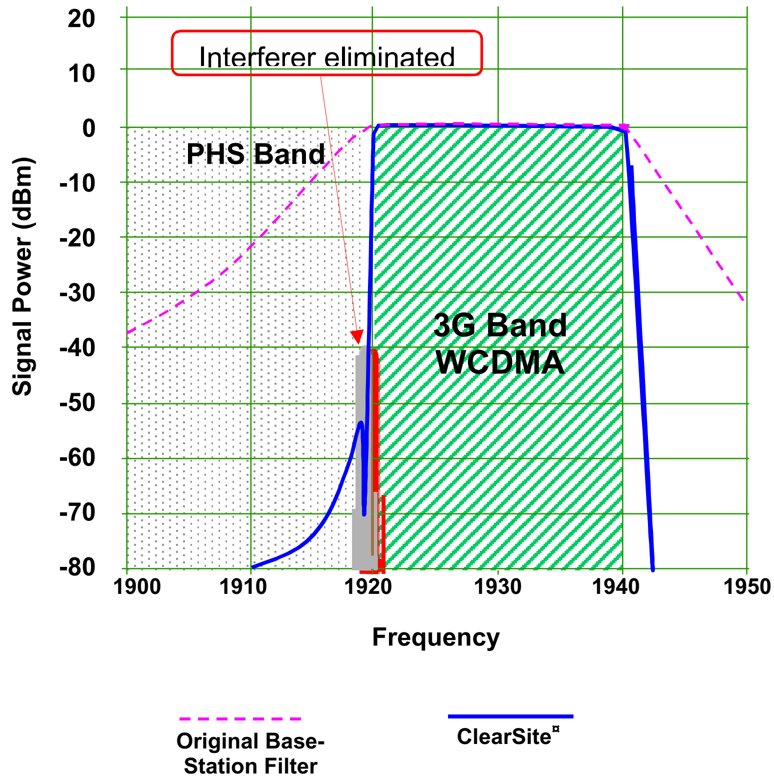

unit, operators have had to become comfortable with issues such as cooler reliability, the time required to achieve operational temperature, and the system response to various failure modes or degradation of performance. As more systems have been deployed and the track record of the technology has been established in the initial deployments, these issues have gradually diminished as major concerns for advanced operators.

Superconductor front-end systems are costlier than the conventional technology that they replace, but the economics of deploying them are increasingly favorable. When capital budgets were large, the answer to problems of coverage, capacity, and QoS was often to build more cell sites. In today's constrained economic environment, alternative solutions are more attractive. As an alternative to adding new cell sites, superconductor front ends are extremely cost-effective, with average selling prices being less than a tenth of the price of a new base station. Even when adding base stations is not a good alternative, but the technology is employed to combat the effects of interference, enhance QoS, and increase data rates, the economics are still very favorable. Based upon the increases in MOU and increased customer satisfaction, the payback period for the equipment can be very short—less than one or two years in many cases.

Superconductor front-end technology is an evolving technology in which new features and enhanced performance continue to be added to the products. The systems are getting more sophisticated, in many instances embodying capabilities for network optimization that are not available with conventional front ends. Some HTS systems are now combined with other base-station components such as duplexers, conventional front ends, and others. A goal of the superconductor industry is to eventually have the technology designed into the base station, which would be the path to the most widespread deployment. Superconductor filter technology itself continues to evolve. Recently, Conductus announced the development of a new filter architecture that provided the most selective wireless filter ever demonstrated—the equivalent of 50 poles. Such filters should be capable of countering the effects of interference in even the most demanding digital wireless networks.

In little more than a decade, high-temperature superconductors have gone from a novel discovery in the laboratory to a technology on the verge of being a mainstay in the wireless industry. Several companies offer a growing set of products that offer unique benefits to increasingly demanding wireless networks. In an age when multiple electronic technologies play a dominant role in our daily lives, superconductors are now joining their ranks.

This article was previously published in the November/December 2001 issue of Base Station Earth Station.

Making DWDM Ready for Prime-Time

The Internet Is Here to Stay, and DWDM Is the Technology of Choice for Alleviating the Bandwidth Bottleneck

Moris Simson
President and Chief Executive Officer
Optenia, Inc.

Introduction

The rapid expansion of the Internet in recent years has inspired a series of new multimedia applications that require high-speed connectivity. These applications continue to demonstrate that the full utility of the Internet is far from having been delivered. E-mail was just the beginning; there are dozens of new "e-things" waiting to be developed into viable business models, all of which rely on the worldwide high-speed connectivity to the Internet. The latter requires broadband access and transmission networks that point to an exponential growth in demand for bandwidth as new users come on stream.

The unprecedented bandwidth demands have required the installation of extensive fiber-optic networks that connect major cities and geographical areas. During the last five years, there has been an increasing reliance on dense wavelength division multiplexing (DWDM) technology to multiply by 40 times or more the amount of information that can be carried on a single fiber. DWDM, which originally was used to provide capacity relief by reducing the need to install more fiber between major urban centers, has started to find new applications. Increasingly, the focus is to reach metropolitan centers and intra-city clusters with high-capacity optical links without having to add additional fiber cables, which can cost up to $60,000/km depending on the dig. Clearly, DWDM ring architectures with optical add/drop multiplexers (OADMs) hold the key to switched connectivity essential to support bandwidth needs in both metro and access networks. These will then feed large office buildings expecting to embrace Gigabit Ethernet and other bandwidth-hungry applications rapidly.

In parallel with this evolution, at the consumer level, broadband access, which today only reaches 3 percent of users, is expected to reach a global penetration of more than 10 percent in less than four years. Never in history has there been such a rapid build-up in bandwidth demand. The key to the surge will be the extremities of the networks or the access points: The more users come on stream and require bidirectional connectivity, the more the bandwidth demands on the backbones will escalate. These broadband applications, and the optical networks required to support them, will drive the knowledge and information-based economy and the way that people communicate and do business. As this transformation unfolds, though, an important question that needs an answer is that of economic viability—Will all of this infrastructure be affordable to build?

The answer is unequivocal: The only technology capable of economical broadband delivery is photonics. The next step in datacom is clearly Gigabit Ethernet for the last mile, representing the link between high-speed optical networks and the individual subscribers. And, DWDM is the technology of choice for transporting and distributing such bandwidth economically.

DWDM

DWDM allows 40 wavelengths or more to be carried across a single fiber. The data from several sources is combined and transmitted optically by a laser of very specific wavelength. Several of these optical signals, each of a slightly different wavelength, can then be multiplexed and transmitted on a common fiber. At the opposite end, a demultiplexer splits these optical signals into individual detectors.

The technology used for multiplexing and demultiplexing narrowly spaced wavelengths has become a critical factor in the cost, and subsequent deployment, of optical networks. Traditionally, this has been accomplished using thin-film filters (TFFs). The assembly of these components is very labor intensive and extremely difficult to automate. Furthermore, TFFs do not offer acceptable performance for any more than 16 channels. Modern DWDM systems are now being designed with 80, 100, or even more channels, requiring much more economical and scalable technologies.

In recent years, there has been a significant interest (and over-abundance of investment) in arrayed waveguide grating (AWG) technology, shown in *Figure 1*. AWGs are manufactured using standard silicon wafer manufacturing techniques and offer better economics and performance than TFFs for more than 16 channels, due in part to the optical chip, which allows for automated assembly. However, beyond 40 channels, the manufacturing of AWGs becomes quite difficult. The chip size, or "die" size, of the AWG is

FIGURE 1

Basic Layout and Operation of an AWG Demultiplexer

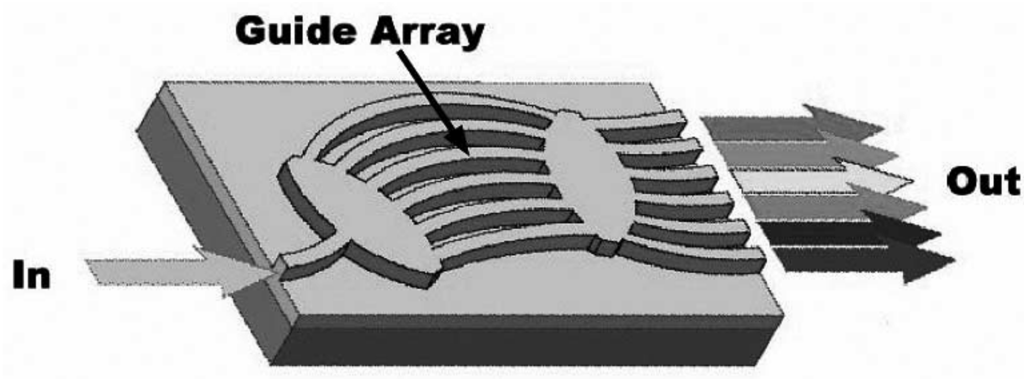

quite large, and wafer non-uniformities result in extremely low yields and reduced performance. The economics are also somewhat compromised by the fact that maintaining the AWG at the required operating temperature requires significant amounts of power, which impacts operational costs and packing densities in the optical system.

The best means of reducing system cost is to integrate more component functionality into one module, which reduces the system complexity as well as the power requirements. The interest in planar waveguide technologies, such as AWGs, stems from the notion that increased functionality could be integrated on the same chip as the demultiplexer. However, with AWGs, the die size is already so large, and yields already so low, that integrating more components on-chip destroys the economic benefits that originally inspired the exercise. A new technology platform scalable to more than 100 channels is required and needs to offer the potential for integration, a very small form-factor, and economics superior to all competing alternatives.

Echelle Grating Technology

Several years ago, there was an acceptance of AWG technology for high channel-count DWDM applications. At the time, the Echelle Grating (EG) approach was hindered by a requirement for very advanced processing techniques, especially in etching. The depth of the etch coupled with the requirements for verticality and smoothness made the fabrication of EGs very challenging. However, with recent advances in semiconductor manufacturing technology, the EG approach now outperforms AWGs and also offers scalability and manufacturability beyond what is attainable using existing technologies. The competitive advantages of the EG include performance, cost, power consumption, size, and scalability.

Although many similarities exist between the operating principles of EG and AWG devices, the layout of each is very different. Both rely on a high-quality grating to separate the narrowly spaced wavelengths. In the AWG, this is accomplished by the waveguide array; however, that section of the device becomes extremely difficult to manufacture at high channel counts or narrow channel spacings. Contrastingly, the EG relies on a traditional reflective grating, created by etching vertical facets into the waveguide core, shown in *Figure 2*. Scaling to high channel counts simply requires adding more "teeth" to that grating, and performance is not sacrificed. The die size, because of the reflective technology, is reduced by nearly a factor of four. That also benefits the power requirements, since the area requiring temperature control is much

FIGURE 2

Basic Layout and Operation of an Echelle Grating Demultiplexer

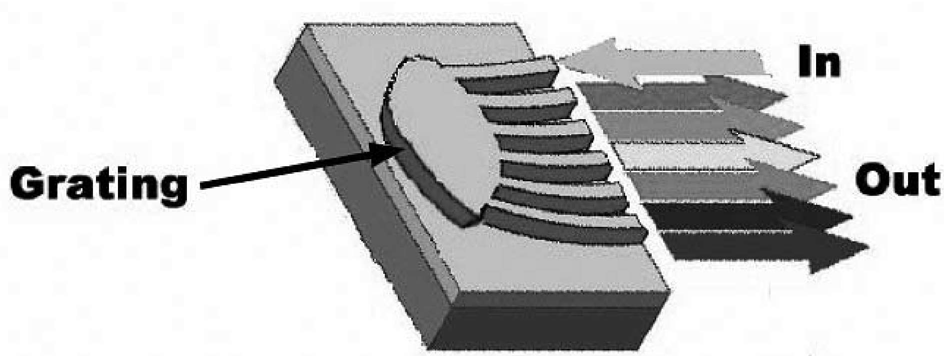

smaller. A typical 40 channel EG requires approximately one fifth the power of a comparable AWG.

For EG devices, the verticality and smoothness of the deeply etched grating facets are critical issues. An example of the type of etch required is shown in *Figure 3*. The reliability and reproducibility of the fabrication process for vertical facets continues to serve as the main challenge for those companies developing and commercializing EG technology. In recent years, these manufacturing issues have been overcome by some companies, and further developments, such as on-chip polarization compensation, have extended the EGs manufacturability beyond that of AWGs.

New Technology, New Opportunities

To be successful, integrated devices must provide greatly reduced cost per function without sacrificing the performance necessary for demanding telecom applications. The difference in the size between AWG and EG devices is illustrated in *Figures 4* and *5*, which compare six-inch wafers containing 40-channel demultiplexers. While only seven AWG DEMUXes can be patterned on a wafer, a total of 36 EG die can fit easily on the same wafer, representing a six-fold improvement in the die-level economics. Because of the small die size, the EG is a suitable platform for the on-chip integration of higher-level functionality.

The small size of the EG device transfers directly into a small package, a feature that is becoming increasingly important for today's complex optical networks that pack more and more capacity into the same size system. A fully packaged 40-channel demultiplexer with an integrated temperature controller is illustrated in *Figure 6*. This is roughly one quarter the size of most packaged AWGs *without* temperature controllers. The smaller EG form factor allows manufacturers to integrate more functionality on-chip or to hybridly integrate the device with an InGaAs detector array to produce very small optical channel monitors suitable for embedding directly in existing network elements. This type of functionality and integration has long been seen as the future of optical components; however, the execution of it has been hampered by the large size and low yields associated with AWGs.

The high performance of EGs results from the small die size, which minimizes the optical loss. The crosstalk is also inherently good because of the quality with which the grating can be defined (see *Figure 7*). Polarization sensitivity is minimized using proprietary compensation techniques, which remain the key to producing commercial EGs of the highest quality.

Die Size: 18 x 20 mm; Adjacent Channel Crosstalk: 35dB; Background Channel Crosstalk: 37dB

Scalability to high channel counts has led many systems manufacturers to shift from thin-film filters to planar waveguide components. However, modern DWDM systems require a large network fan-out to economically deliver

FIGURE 3

Scanning Electron Micrograph of EG Facets Manufactured Using a Reactive Ion Etch

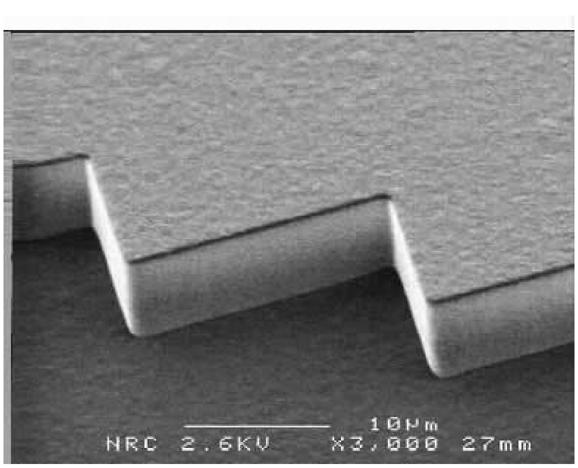

FIGURE 4

40-Channel Silica-on-Silicon AWG Demultiplexers

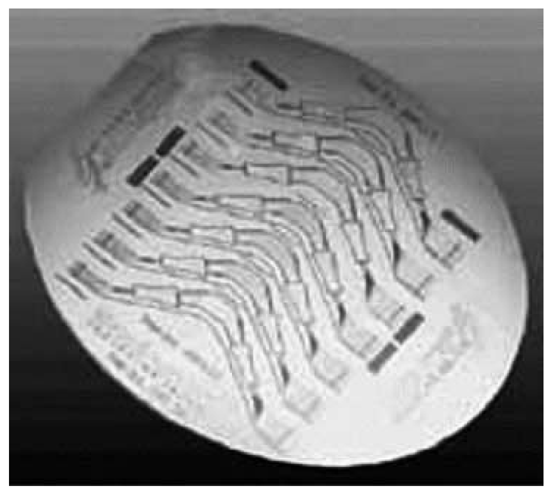

FIGURE 5

40-Channel Silica-on-Silicon EG Demultiplexers

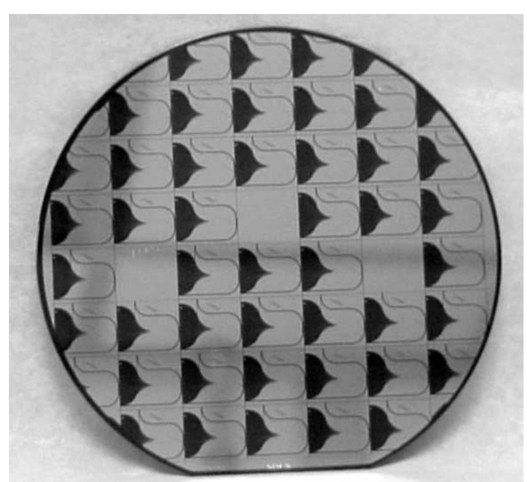

FIGURE 6

40-Channel EG Demultiplexer, Packaged with Integrated Temperature Controller

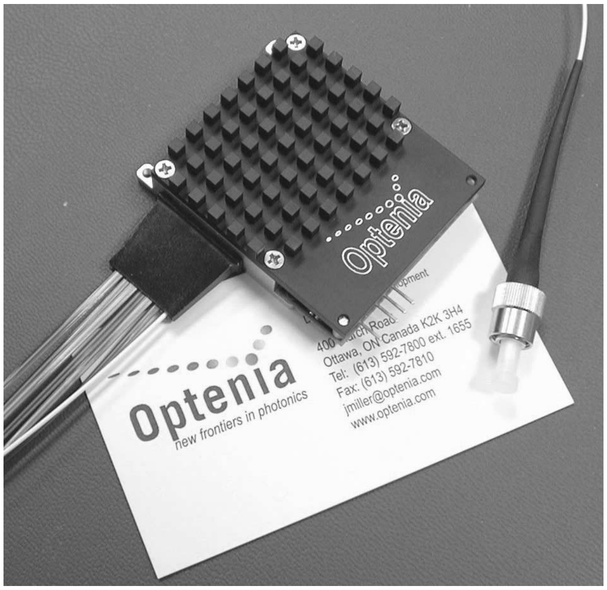

bandwidth to the edges of the network, and that requires scalability not available from TFFs or AWGs. EGs, in contrast, have recently been demonstrated with more than 250 channels on a single chip, and still over a dozen of these chips fit on a standard silicon wafer.

The performance of planar waveguide components, and the resulting signal quality in the optical network, relies on maintaining thermal stability to ensure optimum channel tuning. Both AWGs and EGs are packaged with thermoelectric coolers or heaters to control the devices temperature, but EG devices typically require one-fifth the power of AWGs. This translates directly into operational cost savings, and closer packing densities, which helps reduce the size of modern systems.

Photonic ASICs

Perhaps the most important benefit of EG technology is that it offers the ideal platform on which further functionality can be patterned on-chip. The restrictive size of other demultiplexer technologies makes this difficult, especially at high channel counts, where yields are already so low that adding further functionality results in significantly higher manufacturing costs. EG technology is being coupled with attenuator and switching functions to develop a new form of planar lightwave circuit: the readily customizable Photonic Application-Specific Integrated Circuit (ASIC).
In its most simple form, the Photonic ASIC would consist of a custom demultiplexer design that could be integrated with hybrid detector arrays to develop devices such as embeddable optical channel monitors. Variable optical attenuator (VOA) and switching technology will bring true customizability to this platform and be used to develop fully configurable optical add/drop modules (COADMs) and customer-specific designs for planar lightwave circuits. These modules significantly reduce the complexity and cost of the systems in which they are deployed and are truly the enabling technologies required for tomorrow's DWDM networks.

Conclusion

As the Internet expands and bandwidth requirements continue to grow, DWDM optical networks are being designed to transport staggering amounts of information. These large "pipes" will be fed by smaller metropolitan networks and access nodes, which require high-granularity fan-outs to effectively aggregate or disassemble that information.

As these networks evolve to higher channel counts at lower costs, the ability of today's TFF and AWG technologies to meet future system demands in terms of performance, cost, and manufacturability becomes questionable. The EG approach has quietly matured to a stage where, coupled with progress in silicon-based materials, it outperforms other competing technologies in nearly all respects. It does so at a fraction of the cost of competing technologies while using a fraction of the power. With scalability of up to 160 channels or more, in high volumes and small package sizes, EG technology is revolutionizing planar lightwave circuits and will be a critical enabler of tomorrow's high-functionality Photonic ASIC.

FIGURE 7

Insertion-Loss Characteristics of a 49-Channel EG DEMUX

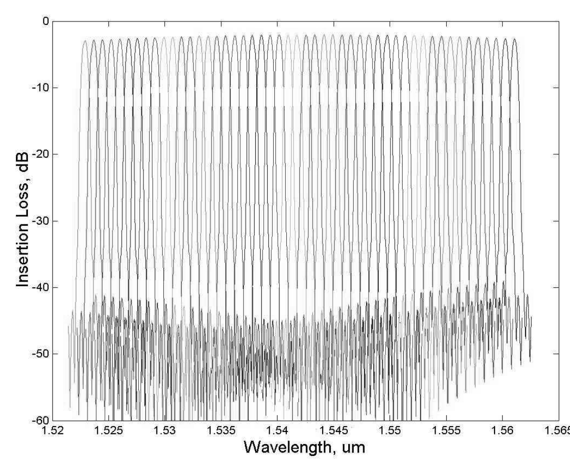

Telco versus IP Service Creation and the Role of Softswitch

R. Paul Singh
Co-Founder and Vice President, Business Development
NexVerse

The outlook in the industry today is quite different from the outlook three months ago, when capital was readily available; every interexchange carrier (IXC) and incumbent local-exchange carrier (ILEC) was going to be under tremendous competitive pressure, and every Internet business model, including the new service-provider model, was expected to generate trillions and trillions of dollars. The recent NASDAQ downward slide on Wall Street has brought back some reality in terms of going back to the business fundamentals; profits (not only revenue) and paying customers (not only subscribers) are needed.

This paper will consider how the industry will evolve to meet the need for service providers to make money from the services they can offer. It will compare the service-creation models of today and tomorrow, the network and services of today and tomorrow, and the coexistence of the old and new models.

Service Creation Today

The basic tenets of service creation, whether based on Internet protocol (IP) or telco models, are the same. Service-creation tools, whether graphical or text based, are used to create services. Once created, services need to be verified and tested and then provisioned for the customers and billed for to be profitable. There is talk of services being created and provisioned by end users—something that will happen but that is not enabled today. However, customizing services through a Web browser with a limited set of predefined choices is already starting to happen.

The following steps apply to any basic service model:

- Service creation
- Service animation and validation
- Support-service deployments and customer scheduling
- Service provisioning
- Hooks to back-office systems (billing, provisioning, and network management)

The Network of Tomorrow

As everyone acknowledges, tomorrow's network will use pure IP or asynchronous transfer mode (ATM) or a mixture thereof. The bottom line is that people will get services off a data network. The intelligence will be distributed to the edges as well as the core (see *Figure 1*). At the corporate edges will be corporate gateways (also known as integrated access devices [IAD]) that will connect the private branch exchanges (PBX) and local-area network (LAN) traffic on one integrated data connection. If the corporation does not have the PBX and/or voice-mail systems, the voice services will be provided by an IP Centrex–type server off the network core while unified messaging may be provided by another server in the IP cloud. The same process is happening in the home with residential gateways or IADs proliferating with broadband cable and digital subscriber line (DSL) connections, wherein the personal computers (PC) are connected to the Ethernet port while black phones are connected to the analog ports on the gateway. Different application servers in the cloud offer different voice, data messaging, and streaming services. Then there are mobile IP phones and personal digital assistant (PDA) users who will also use the services enabled by this next-generation network (NGN).

In this new architecture, softswitches are needed in the network not only to provide the network intelligence to these edge devices so that calls can be sent to the appropriate devices but also to be able to terminate calls back to the existing billions of black phones connected directly to the time division multiplex (TDM) side of the public switched telephone network (PSTN). Softswitch architecture, done right, will enable the creation of true application servers in the telecom network and make the communications application service provider (ASP) model viable. Softswitch hides all the connectivity nuisances, be it with a signaling system 7 (SS7)–based network, or the primary rate interface (PRI) or T1 communicating applications specification (CAS), or simply via an IP/ATM access device and present industry-standard open interfaces to the application server, much like an operating system (OS) in the data world hides the nuisance of PCs, servers, and other peripherals from the applications.

There is a school of thought in the industry that everything will become end-to-end IP and replace the PSTN. The industry is not there yet. It cannot change to tomorrow's

FIGURE 1

Tomorrow's Network: Integrated Packet Network for Voice and Data

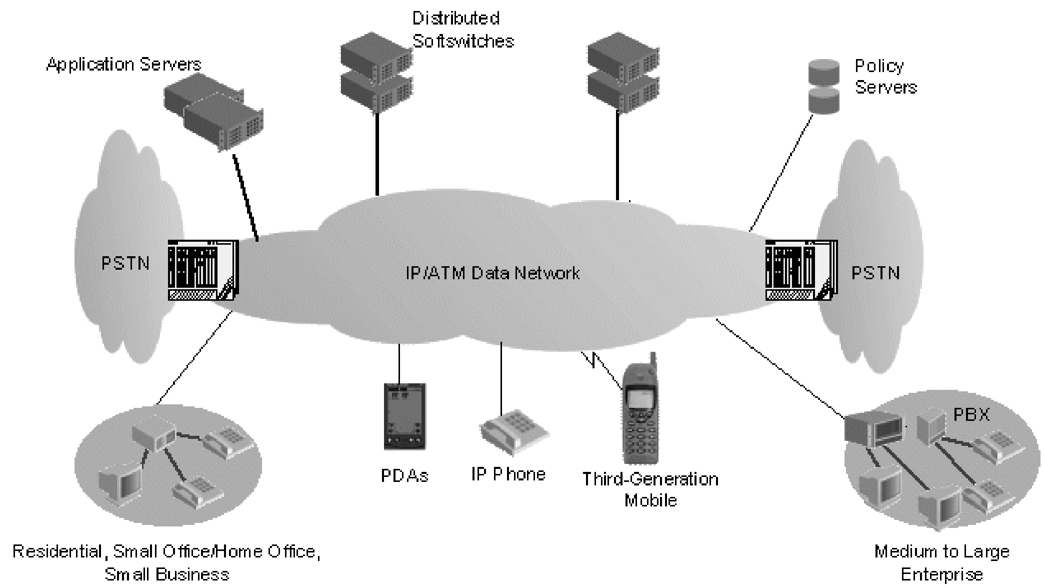

network with the flip of a switch because there is a hundred-year history with the PSTN. The PSTN will not be discarded because it is already established and it still works. Therefore, the gateways that bridge calls between PSTN and IP networks will continue to exist. There may not even be any reason to replace the good old black phone for a grandma who has only $2 on her phone bill every month. Instead, service providers target the digitally empowered users who have broadband access, PDAs, and cell phones; who are the early adopters of the new technology; who spend more than $100 per month on communication; and who want customization of the services. However, unless the new network achieves the same reliability as the PSTN, there will never be a complete transition into the new network.

In summary, tomorrow's network will be data based, highly scalable, and open. The new network will not appear in one day but will evolve over time, through transitional steps. Some predict it will take two to five years; others predict 10 to 20 years. As the services and networks are built, it is important to make sure that there will be a way for service providers to make money with them, because only then will new equipment be deployed. No new services will ever be created without a revenue model that satisfies the service providers that must deploy the equipment.

The Emerging Service Models

User-, Network-, and Application-Centric Providers

The current model of service providers in the telecom is born out of the Bell system, in which ILECs own the network, the customer, and the application. On the data communications side, this model did not work, as evidenced by the separation into the following three different kinds of service providers:

- User-centric service providers, such as America Online (AOL)
- ASPs, such as Coreo
- Network service providers (NSP), such as Sprint, Level 3, and others, that sell to such companies as AOL

Telecom service providers, over time, will also split into these segments. There is some evidence of that in the mobile space where Virgin Mobile, a user-centric service provider, is buying the network services from others. ASPs will offer specialized vertical applications that focus on particular niches. The user-centric service providers will buy services from the NSPs and ASPs and create a revenue model around that strategy.

Vertical Services

The world is moving away from the "one size fits all" types of services that are offered by phone companies today. In the data world, users have a lot of control and capability to customize their services, and there are many vertical services—for example, for the manufacturing and services industry or for a doctor's or an attorney's office. In the converged world of voice and data, the users will demand the same level of control and customization, and service providers that offer these vertically integrated services will succeed and achieve faster profitability.

Also, the industry is going to move from a service-provider-control model to an end-user-control model, but it is unclear how much control will be transferred and when. However, service providers that offer levels of customization ahead of their competition will capture this highly profitable but demanding segment of digitally empowered users. In addition, there is a need for vertical services focused on individual professions, such as doctors, attorneys, and salespeople. So service providers that want to increase their profitability

FIGURE 2

Next-Generation Architecture and Interfaces

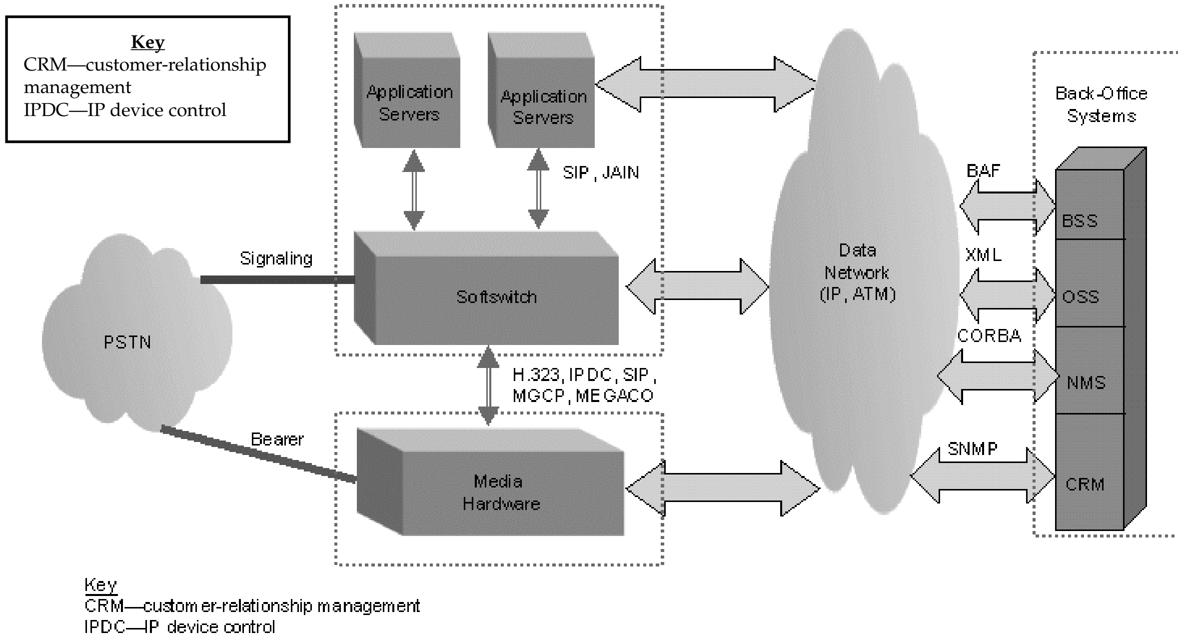

and achieve differentiation will need the new service-creation model and infrastructure technology that enables new levels of control and customization.

Billing Reforms

Today, a typical consumer deals with four different service providers: one for Internet access, one for local telephony, one for long distance, and yet another one for Internet access. Yes, people want choices, but they would prefer to write one check for all the services that are integrated into one bill. The billing models for wireline telephony will also undergo significant changes. The classical example is the cell phone, where the user pays for a set number of minutes and normally stays within that number, or the new long-distance models with phone service for seven cents per minute. Removing the billing effort that is spent by the existing phone companies and implementing a simplified billing structure would generate significant savings, exactly as a flat income tax would reduce administrative costs. Everybody would benefit, and economies of scale would be achieved. The bottom line is that the way services are billed and sold will undergo major changes in the next few years.

The Next-Generation Architecture and Interfaces

It is clear that the next-generation architecture will be a decomposed model, will be multivendor and multiprotocol, and will have industry-standard interfaces between these best-of-breed systems (see *Figure 2*).

Decomposed Model

The next-generation converged services will depend on a decomposed architecture for achieving flexibility. Exactly as in the data communications world, in which the mainframe was split into a distributed architecture, the "telecom mainframes"—i.e., the circuit switches—are undergoing the same transition to a three-layer model, in which the bottom layer is the hardware, the middle layer is the softswitch, and the top layer is the applications. Much like PCs and peripherals, the various types of devices that form the hardware layer in this new architecture include remote access servers (RAS), media gateways (MG), IADs, IP phones, and media servers.

The middle layer, the softswitch, is much like the OS and it controls the hardware and provides open interfaces to applications. With a scope as wide as this, it is clear that the new infrastructure will be multivendor.

Multiprotocol

The next-generation architecture, much like the current systems, will be multiprotocol. As seen in data communications, old protocols, once they work, continue to hang around the network while the new ones enter into the system. So it is important that the softswitch support multiple protocols to enable new and existing services to coexist. It is also important for the softswitch to be able to shield applications from the nuisances of different hardware and protocols they may be running.

Interfaces

There are many interfaces between these layers that are in different stages of standardization. PSTN interfaces—including SS7, PRI, CAS, and many others—need to be supported by both the hardware and softswitch. The International Softswitch Consortium (ISC), the Internet Engineering Task Force (IETF), and the International Telecommunications Union (ITU) are driving many of the new standards in the industry. Media interface protocol

H.323 still exists in many networks, while session initiation protocol (SIP) and MG control protocol (MGCP) are gaining ground. MG control (MEGACO) is right behind and will start to enter the networks in the next 12 to 18 months. As with any protocol, many of the interfaces are defined, but many of the implementation profiles exist and many new ones are being added to ensure smoother interoperability between different vendors.

The various application interfaces between softswitch and applications include SIP, Java application programming interfaces for integrated networks (JAIN), Parlay, and extensible markup language (XML). The ISC is working to make sure that the SIP interfaces between the application servers and the softswitches are defined. Many of the existing applications that use H.323 can also be interfaced by using H.323 standards.

In addition, the whole system has to be billed for, provisioned, and managed, and there are many interfaces that exist between softswitch and back-office systems that make up the business support system (BSS), operations support system (OSS), and network-management system (NMS) applications. In the billing area, the Bellcore automatic message accounting format (BAF) interface continues to be king, while Internet call detail record (ICDR) and other efforts are under way and will institute billing reforms. In the management area, simple network management protocol (SNMP) is the dominant system, although XML and common object request broker architecture (CORBA) are popular for integration with provisioning systems.

Next-Generation Services

Much talk about the new architecture is based on the premise that the applications will be completely IP based, providing services from anywhere to anywhere. Many communications ASPs have emerged in the market and were forced to build communications infrastructure because the next-generation architecture has not yet been deployed. Once deployed, the Web-hosting model will apply to the communications ASP, which will be able to offer services from anywhere to anyone with new IP communications applications.

Many of the applications that will be deployed first will IP–enable existing applications. For example, voice mail and e-mail combined with unified messaging will be deployed on the next-generation architecture on IP backbones such that users can get services from anywhere. Similarly, Centrex services will become IP enabled and hence will be offered not only by local-exchange carriers (LEC) but also by some service providers that will use a hosting model of Centrex services. The application list includes unified messaging, conferencing, unified call centers, communication portals, presence-based servers, and IP Centrex.

The New Service Model

The core services being built are very similar to what was on their old networks, but the new services are much more flexible and are based on a packet network. People are opening up to many different services and are enabling new service models to exist.

Three popular service-creation models are being discussed in the industry, and they each will play a particular role because their unique strengths and weaknesses fit with the service provider's business models. The three models are as follows:

- Edge-based service creation
- XML–based service creation
- JAIN/Parlay–type service creation

In the edge-based service-creation model, it is assumed that every application server will have its own service-creation model. It will provide a service-execution environment and will call upon the services from the softswitch as needed, using such protocols as SIP. This model is somewhat similar to the data communications model. However, much work is needed to unify these services and to offer a revenue model to the service providers.

In the XML–based service-creation model, the services may be created in an off-line world but will run on the softswitch or on some component of the softswitch. It may be more tightly integrated with the softswitch than the last model but will make it easier to manage and control the services.

In the JAIN/Parlay model, it is the interfacing with the application programming interfaces (API) supplied by these models that will be provided by the softswitch.

Which of these models is the best and will prevail? All these models will coexist, although one model will be better suited for one type of application/service, while another model will be used for a different service. Therefore, it is important that the softswitch be designed correctly to accommodate these different models.

Summary

In summary, packet networks are the future, but the transition to packet networks from TDM networks is an evolution—not a revolution. To make this transition smooth, an open architecture is needed with focus on software and applications with flexible service-creation environments (SCE). The network architecture will be decomposed, with vendors choosing their core strength in providing media hardware, the softswitch, or the application servers.

The applications deployed today in this architecture are the ones that are leveraging or offloading some of the existing switches. However, the growth of this new NGN architecture will occur only when there are broadband systems and new applications that existing circuit switches cannot create. It is a challenge for the whole industry, including the service providers and the vendors, to make the multivendor systems work, because failure would revert the industry to a single-vendor, proprietary system.

Second-Generation, Multi-Mode Wireless LAN Clients

Yoram Solomon

Vice President and General Manager, Advanced Communications Product Group

PCTEL, Inc.

Abstract

In this article, I am addressing the future of wireless local-area network (LAN) connectivity, referring to it as the second-generation (2G) wireless LAN. The four major issues that need to be addressed by wireless LAN product manufacturers are 1) the coexistence of physical-layer standards such as 802.11a, 802.11b, and 802.11g, 2) security, quality of service (QoS), and other "maturing" factors, 3) interference immunity, and 4) the user experience.

For every item, I am describing the current "state of the industry," as well as what will need to take place in order to have mature products ready for mass deployment. This article starts and concludes with a vision of the proliferation and mass-adoption of wireless LAN through the deployment of the 2G wireless LAN products beyond today's "early adopter" market.

Foreword

In the first quarter of 2001, the personal computer (PC) market continued its slowdown, as did the entire economy; however, the revenue from wireless LAN products was reported to grow 15 percent quarter over quarter. Has this market matured?

Today's wireless LAN market is in its "early adopter" stage. Today's products are not as simple to install and not as seamless to operate. The users of wireless LAN products are typically very technical and can overcome the hurdles associated with the installation and use of wireless LAN. However, the benefits of using wireless LAN are tremendous. Personally, I use only wireless LAN at the office, roaming from office to conference room. At home, I get connected to the Internet through my cable modem and wireless residential gateway, using my computer anywhere in the house. When I travel, I use my MobileStar subscription to connect to the Internet at booths located near almost every American Airlines' gate and make the most out of the downtime before my flight. Oh, and by the way, I also use it in the Starbucks near my house while I sip my "*grande coffee Frappuccino*®."

There is no doubt that wireless LAN is going to be adopted. However, there are several trends that must be observed and acted upon in order to deliver the right products for mass adoption. This article covers the four most important trends and issues around wireless networking. Addressing those correctly will promote wireless LAN adoption beyond our wildest dreams (Did anyone ever believe that there would be 400 million cellular telephones sold in one year?), introducing the 2G wireless LAN.

Frequency Bands, Modulations, 802.11g, and the PHY Layer

Wireless LAN products can use several frequency bands allocated by the Federal Communications Commission (FCC) and other regulatory agencies worldwide. The most common frequency bands are the Industrial-Scientific-Medical (ISM) frequencies of 902–928 megahertz (MHz) and 2.4–2.4835 (gigahertz) GHz, and the Unlicensed National Information Infrastructure (U–NII) frequencies of 5.15–5.25 GHz, 5.25–5.35 GHz, and 5.725–5.825 GHz.

Due to the lower bandwidth available in the 900 MHz band (only 26 MHz), the achievable data rates in that band are relatively lower, and with the increased use of cordless telephones in that band, it has never been widely used for high–data rate wireless networking.

The 2.4 GHz band was the first one adopted widely for wireless networking, mainly by the 802.11b standard (also known as Wi-Fi), by the wireless home networking standard (HomeRF), and by the wireless personal-area networking standard (Bluetooth), known as PAN. One of the leading reasons for such a wide proliferation of products and standards in this band is the availability and feasibility of silicon operating components. The resulting cost reductions, mostly through moving to low-cost complementary metal oxide semiconductor (CMOS) processes, increased the penetration and deployment of such products.

The 5.4 GHz band is not yet used for such products, as the availability and feasibility of silicon components in these frequencies are still being evaluated. The Institute of

Electrical and Electronics Engineers (IEEE) created a standard known as 802.11a that uses the 5.4 GHz band for high-speed networking, and several companies are developing components to support this standard and offer data rates higher than Wi-Fi in this band. It is expected that the first products will hit the market toward the end of 2001, followed soon by price reductions, causing this standard to then be widely adopted.

While the 802.11b standard offers a data rate of 11 megabits per second (Mbps) at the 2.4 GHz frequency band, the 802.11a standard offers a data rate of 54 Mbps at the 5.4 GHz band. Much like the migration from 10 Mbps to 100 Mbps (and then to Gigabit Ethernet and beyond) for wired LAN, there is a perceived need for higher data rates in the wireless LAN arena, beyond the 11 Mbps offered through Wi-Fi products. Texas Instruments (by acquiring Alantro) has offered a new modulation scheme called packet binary convolutional coding (PBCC), capable of delivering data rates of up to 22 Mbps still using the 2.4 GHz band. This technology was proposed to the IEEE for a working group called 802.11g. Shortly thereafter, Intersil offered a counter-proposal using orthogonal frequency division multiplexing (OFDM) capable of delivering 36 Mbps in the same frequency band. Initially, both proposals violated the FCC rules (15.247) for the use of that frequency band. However, the new FCC proposed ruling (Docket 99-231, discussed later in this paper) suggests modifying the rules such that neither proposal will be in violation. Meanwhile, in the IEEE 802.11g working group, the TI proposal was voted off, but the OFDM proposal still needs to get the support of 75 percent of the members to become the 802.11g standard.

What will the impact of 802.11g have on the use of 802.11b and 802.11a? The effect on the current 802.11b (Wi-Fi) standard is pretty predictable. Once a higher-data rate standard gets approved, the existing 802.11b will simply become a fallback data rate to the higher–data rate standard. This is similar to the original 802.11 (supporting 1 Mbps and 2 Mbps) that became the fallback data rate to the 802.11b (driving 11 Mbps or 5.5 Mbps). The main question is—What will the impact of a higher–data rate 2.4 GHz standard (802.11g) be on the higher frequency, higher–data rate 802.11a standard? I will assume two scenarios, depending on the first standard to hit the market with products.

Scenario 1: 802.11g Products Available before 802.11a Products

With the OFDM modulation for 802.11g being able to deliver as high as 36 Mbps, the main question is whether 802.11a will still be needed. I believe that there will still be a need for 802.11a for several reasons. One reason is the allocation of additional bandwidth that provides a higher aggregate data rate usable for wireless LAN. If 802.11g is used within a certain physical area, then only 36 Mbps becomes available in this area. If, however, both 802.11g and 802.11a services (or coverage) are offered in an area, then the aggregate data throughput offered at this area increases to 90 Mbps and therefore can support more users. Furthermore, the 2.4 GHz band sees more interference today from non-wireless LAN products, such as Bluetooth devices and cordless telephones. Adding the 5.4 GHz band for wireless LAN applications will help in avoiding some of that interference. The conclusion from this scenario is that the introduction of 802.11g will not delay or prevent the later deployment of the 802.11a-based products.

Scenario 2: 802.11g Products Available after 802.11a Products

As in Scenario 1, the introduction of 802.11a products before 802.11g products will not prevent 802.11g products from being deployed. While the 5.4 GHz band is limited to 54 Mbps, having 802.11g deployed will increase the aggregate data throughput by 36 Mbps to the total of 90 Mbps. Furthermore, with the 802.11a products suffering from a shorter range, there may be somewhat different applications for both standards. In any event, the conclusion remains the same: The deployment of 802.11a products before that of 802.11g products will not delay or prevent the later.

Given these conclusions, it seems that there is opportunity for the 5.4 GHz and 2.4 GHz standards to coexist.

Multi-Mode Access Points

The next question that comes to mind is—What will the migration path be for deploying the high–data rate standards?

One possibility is that all existing (installed) access points will either be upgraded to the high–data rate standard or become multi-mode. Most of the installed access points are not scalable or upgradeable to the new standards or do not support multiple standards. Therefore, the scope of this proposition is that the access points will be replaced in their entirety. This is a very expensive proposition and one that will probably never happen, much like the old analog cellular infrastructure was not replaced once personal communications services (PCS) and code division multiple access (CDMA) were introduced. I believe that we will continue to see multiple 802.11b access points, along with new multi-mode access points.

Another question is whether there will be new access points that support only the higher–data rate standards but not the existing 802.11b standard. With the current rate of 802.11b client product deployment, it has been estimated that by mid-2002 (when we expect to see some 802.11a or 802.11g products available), there will be some 20 million 802.11b-only client adapters installed in computers and other devices. Installing access points that do not support this 802.11b standard means that 20 million users will not be able to access them, and that is unimaginable.

Multi-Mode Client Devices

Now comes the question of whether the new clients are going to support only a high–data rate standard. Assuming that not all of the currently installed (802.11b) access points will be replaced by multistandard or high–data rate standards due to the required investment in infrastructure, there will still be many 802.11b-only access points installed. Therefore, installing wireless network cards in computers that support only the high–data rate standards would disable these computers from accessing the existing 802.11b access points. Wireless Internet service providers (W–ISPs)—such as MobileStar, WayPort, AirWave, and others—have already installed access points in public places such as airports, hotels, and even Starbucks' locations. They probably will not replace the existing 802.11b-based infrastructure, and users wanting to access the Internet or their corporate intranets in these locations will have to have equipment supporting the 802.11b standards. MobileStar, for example, offers its users both 802.11b cards and frequency-hopping (FH) cards to support the older infrastructure for the lower–data rate FH standard. However, the problem with that is much less acute than it might be if they had to provide all of their users with both a high–data rate card (supporting 802.11a and/or 802.11g) and an 802.11b card.

As new computers and other non–PC devices (such as personal digital assistants [PDAs]) become smaller, the amount of space allocated for extension accessories (such as wireless LAN cards) becomes more limited. Installing two cards (one to support 802.11b and one to support 802.11a) will be nearly impossible.

Moreover, wireless LAN is becoming an integral part of new PCs, as PC manufacturers such as Dell, Compaq, Toshiba, IBM, and others are shipping new computers with bundled wireless LAN cards in a mini–peripheral component interconnect (PCI) form factor. While PC–card (PCMCIA) accessories can be replaced during operation, the embedded mini–PCI cards cannot.

Therefore, the best solution for multistandard wireless LAN client support is multistandard wireless LAN cards, supporting more than one standard on the same card. These cards should support 802.11b, 802.11g, and 802.11a due to their coexistence in the marketplace.

Service Discovery, and 802.11b as the Negotiation Standard

With the aforementioned coexisting environment, products will need to implement "intelligent" service-discovery procedures, allowing them to identify the available networks and to select the best one to use.

The assumption is that beginning in 2002, there will be a mix of clients, some of which will have an 802.11b-only access card and some of which will have high-end multistandard cards capable of higher data rates. There must be a common way to communicate with them. I have proposed that this would be accomplished through the use of the 802.11b standard at its basic form. Since all client devices will be able to communicate through 802.11b, it can be used as the initial link between the client and establish the access point. Then, still using the 802.11b standard, the client and/or the access point will attempt to negotiate a higher level of communication, either using a higher–data rate standard (such as 802.11a or 802.11g) or enabling a higher level of security or QoS, or other enhanced features. If both ends (the access point and the client) can support a higher data rate or enhanced functionality, they will negotiate their link up to the highest common ground.

Security, QOS, and the Upgradeable MAC Layer

Security

In February 2001, the Internet, Security, Applications, Authentication and Cryptography (ISAAC) research group at the University of California, Berkeley, published an article attacking the level of security provided by the wired equivalence privacy (WEP) protocol supported by the Wi-Fi standard. Following this publication, the public perception of the security offered by wireless networking was that these networks are not to be trusted. A not-very-well-known fact is that most equipment vendors and W–ISPs recommended that WEP be disabled altogether "to simplify the installation and operation" of such equipment.

The IEEE formed a working group named 802.11i to address the security needs of the wireless LAN user community. Two future developments are certain: Some action will be taken, and it will render the existing wireless LAN cards obsolete. These changes will most likely affect the media access control (MAC) layer, which is implemented within the chipset that constitutes the core of the wireless LAN card. When a new security standard for wireless networking emerges, new chipsets will be developed, installed on new cards, and sold to end users.

Quality of Service

The basic 802.11b standard has one disadvantage when compared to HomeRF and Bluetooth: its inability to guarantee QoS. This disadvantage translates into an inability to deliver non-data services (such as voice and video) in a reliable, high-quality manner. Once again, the IEEE addressed the issue by forming a working group called 802.11e that proposed a QoS standard be implemented in the MAC layer. If approved, the next-generation wireless LAN products will support multimedia applications with a high level of quality. Once again, when QoS is implemented, it will render the existing wireless LAN products obsolete. They will be incapable of supporting the new standards and will force the end users to replace their hardware.

The Upgradeable MAC

The introduction of new standards that address QoS, security, and other issues will automatically make existing products obsolete. End users are typically reluctant to purchase products that might become obsolete and prefer to wait until the next generation is available. The way to reduce that risk is to offer a field-upgradeable MAC layer that is fully programmable in real time. When a new standard is released, the MAC can be upgraded by simply downloading a new driver.

This might appear to be a simple task, but the key to its successful execution is to correctly forecast the future standards that will be released and to incorporate a powerful-enough programmable platform that will be capable of supporting these standards. The trade-off can be in terms of flexibility versus cost and flexibility versus power consumption. A higher level of flexibility can be achieved when functions are implemented in software rather than in hardware; however, this typically involves a higher-cost platform (multiprocessor) and consumes more power due to the high-speed processors deployed. Moore's Law and the overall progress in semiconductor development will compensate for these trade-offs in the long run. The correct trend still needs to be a shift from hardware implementation to software implementation, offering the user the programmability and upgradeability that will reduce the risk of purchasing products before they fully mature.

FCC, IEEE 802.15.2, and Interference

On May 10, 2001, the FCC proposed changing the rules for the unlicensed 2.4 GHz ISM band. The commission proposed "to reduce the amount of spectrum that must be used for frequency hopping spread spectrum systems operating in the 2.4 GHz band (2400–2483.5 MHz), and to eliminate the processing gain requirement for direct sequence spread spectrum systems." This new proposed rule is known as "Docket 99-231." This ruling will allow both techniques proposed to the IEEE 802.11g committee (Texas Instruments' PBCC modulation and Intersil's OFDM modulation) to comply with the new FCC rules. This ruling will also allow FH systems (such as Bluetooth and HomeRF) to avoid frequencies used by a Direct Sequence 802.11b (or 802.11g) system, not having to hop throughout the entire frequency band.

While this ruling would seem to promote the adoption and deployment of wireless networking in this frequency band (through higher data rates and better mutual avoidance of the existing standards), it also allows for greater interference in this frequency band.

We must remember that this unlicensed band is open to use by other types of devices, such as cordless telephones. The FCC previously regulated all of the products using this band so that they had less interference with one another by spreading the signal over a wider band (whether through higher Direct Sequence processing gain or through a higher number of frequency hops in a FH system). With the proposed ruling, a device now operating in this band might generate a more "concentrated" interference with another device. Even before the ruling, devices operating in the 2.4 GHz band included cordless telephones, microwave ovens, and others. Now, with the new ruling, the expected interference from cordless telephones will increase.

The IEEE has addressed the mutual interference between Wi-Fi (802.11b) devices and Bluetooth devices through a newly formed working group named 802.15.2. This working group, however, is proposing to solve only the mutual interference created by Bluetooth and Wi-Fi products and does not address any interference to or from other devices using the same band.

Today, the deployment of products in the 2.4 GHz band is relatively low. However, we are expecting wireless LAN (mainly Wi-Fi), PAN (mainly Bluetooth), and even cordless telephones operating in this band (and thus offering longer range and multiple handsets) to proliferate. Once that happens, the level of man-made interference in this band will increase significantly. Devices incapable of avoiding such interference will suffer degradation of quality through lower throughput (or completely loose connection), range decrease, and noise. On the other hand, devices with intelligent interference-avoidance capabilities will outperform their competitors, offering superior range, throughput, and overall quality.

The User Experience

Microsofts announced that the Windows® XP operating system, launched in October 2001, will support wireless networking in a way that no other previous operating system has done before. This support addresses zero configuration, security, authentication, roaming, and service discovery. At the Wireless Ethernet Compatibility Alliance (WECA, the promoter of the Wi-Fi log and certification) annual member meeting in Helsinki in June 2001, Microsoft even promoted the idea of multi-mode wireless LAN client radios to simplify the user experience.

Where are we today? The installation of an 802.11b product is somewhat complicated. A one-hour installation can easily turn into a six-hour installation. Roaming between access points is not trivial, and extending this roaming to cover residential gateways and public "hot spots" (such as hotels, airports, and even Starbucks' locations) requires some "unnatural acts." Having to change configuration, release the current Internet protocol (IP) address, renew the IP address, and other tasks that need to be performed reduces the pleasure of experiencing a wireless LAN. To compare this experience to that of a cellular telephone, imagine what it would be like if, when you roam between cell sites (every 1–2 miles), you had to change your cell-phone configuration, release its telephone number, renew its telephone number, and do all of that so you could still accept phone calls. Unimaginable?

From this comparison, you can understand that in order for wireless LAN to be as widely adopted as cellular telephones, its operation must be as seamless as that of cellular telephones. While Microsoft® XP might offer a solution, the solution should be supported with existing operating systems that are still bundled with new computers.

Conclusion

In this article, I described 2G wireless LAN clients. The coexistence of the different frequency bands used for wireless LAN will lead to the adoption of multistandard (or multimode) products. The proliferation of protocols and functions (such as QoS and security) in an immature standard will lead to the development of software-programmable products that can be upgraded as the standard matures. The ability to avoid the growing in-band interference will differentiate good products from bad ones and give the manufacturers of the better products a competitive advantage. Last, but not least, the user experience will have to be as close as possible to that of a cellular telephone, with seamless operation that can leverage on the pervasiveness of wireless LAN in our lives.

Telecommunications Marketing Best Practices

David J. Sparks
Executive Vice President and Director of Client Services
Robinson & Maites Inc.

Marketing spending has dramatically increased in the telecommunications industry since deregulation in 1996. Spending has been driven by an increase in the number of competitors, greater geographic scope, a vast array of new services and products, and, of course, by open market competition. As a result of a slowing economy and general uncertainty in the consumer and business markets, we are currently experiencing a demand contraction and a reduction in aggressive spending by most major telecommunication competitors. In an effort to help guide telecommunications marketers in what may be a very difficult time ahead, we have reviewed marketing practices across all major firms over the last two years and identified what we consider to be telecommunications marketing "best practices."

Our firm has supported a range of telecommunications companies since the early 1980s. We have either worked with or competed against most of the major wireline, wireless, and data companies in the areas of strategic planning, new product introductions, and customer acquisition and retention. This article is based on research conducted in North America and on our experience and perspective in this rapidly changing marketplace.

Challenges Facing Telecommunications Marketers

Telecommunications marketers face a number of challenges in effectively and efficiently marketing their products and services. We believe these challenges are common throughout the industry:

- Difficulties with deregulation
- Rapid technological advancement/change
- Increased competition for customers
- Ponderous bureaucracy
- Industry consolidation
- Changing customer needs and expectations

Each of these challenges makes marketing a very complicated subject for all involved. In fact, many believe that telecommunications is the most complex marketing environment in the business world today.

The aforementioned challenges have created a common set of problems facing telecommunications marketers. We have utilized these problem areas to frame and identify best practices.

Best Practices: Marketing Solutions for the Most Common Problems Facing the Telecommunications Industry

1. ***"Learning Builds Earnings"***

   Problem: Because of severely limited customer knowledge, telecom customer marketing initiatives tend to rely on the most inefficient marketing tactics, namely mass advertising. Customer knowledge is limited because of regulatory constraints on information sharing and corporate silo infrastructures, which impede enterprise-wide understanding. Telecom marketers bang away at customer households and businesses as if they were single entities. In fact, increased knowledge would indicate that each has multiple individuals or groups with different needs. Discrete knowledge of individuals and internal corporate groups is needed to understand specific needs.

   Best Practice: **Increase customer knowledge continuously through the use of a closed-loop Learning Marketing Model to improve sales, upsell, and retention.** This is accomplished by making every tactical execution a learning experience. Relevant data is collected and reused to refine and enhance marketing efforts to better focus targeting, product offerings, and incentives. This leads very quickly to more effective and efficient marketing execution.

   Examples: The GTE Communications Corporation utilized this approach in launching their competitive local-exchange carrier (CLEC). They applied a Learning Model to all marketing activities, which better directed segmentation, messaging, product emphasis, incentives, and media mix. The Learning Model shifted their targeting efforts to specifically identified high-value customers. Targeting was based on behavioral data so that communications and marketing activity was triggered by calendar events or "moments of truth." GTE was able to use actual customer feedback to model anticipated behavior and/or lifestyle changes.

The results were vastly improved acquisition rates, lower cost of acquisition, reduced talk time at call centers, and dramatically increased close rates. It also enabled this CLEC to improve upsell and reduce churn.

2. "Minimize Customer Risk to Maximize Sales"

Problem: For telecommunication companies, existing products and services are often the most profitable to market to their embedded base. These products and services are available for immediate use and have broad penetration capability with few real costs other than marketing. The problem is that for customers, these products (e.g., vertical call-management services) may be the least appealing because of apathy, lack of benefit, or poor understanding of how they would be used. Because these products and services are not based on real or perceived customer need, telecom marketers are forced to spend heavily to generate sales. A subsequent problem is a high degree of churn as customers realize that they do not really need or understand how to use the product.

Best Practice: **Utilize special high-profile promotional offers to consistently overcome consumer intimidation/lack of understanding and stimulate trial.** Telecom marketers should use special promotional offers to allow customer to "try" at no risk. This reduces the cost of marketing and allows a greater number of people to try, thus enabling self-selection for those who really need the product or service.

Example: To motivate customers to try or buy additional services, Verizon Network Services built a "preferred pricing" model into their bundled array of product offerings. "Big Deal" customers got Verizon's best prices for buying into a very flexible bundled package. Options allowed customers the flexibility to choose other services at Verizon's best rates. By offering a "preferred" pricing model to customers, it reduced their hesitancy to try new products and services. Verizon had double-digit increases in the sale of add-on services and upgrades as a result of this effort.

3. "Less Is More"

Problem: Traditional telecom sales channels, such as call centers, do not maximize sales opportunities. While call centers handle both sales and service, the majority of the calls are for service. Call centers' ability to effectively sell is limited because of training and the high costs to initiate a selling process with inbound calls. When it does happen, product, pricing, and ordering are complicated affairs with multiple offers, system issues, and state variations. This results in low productivity and a less than satisfying customer experience.

Best Practice: **Simplifying call-center channel complexity makes it easier for agents to sell and easier for customers to buy. Products, services, and promotions should be constructed specifically for channel sales.**

Example: GTE Network Services accomplished this by focusing on a single offering that could be effectively supported at the channel level. They ran a vertical services program that allowed customers to choose any vertical service for a trial period of 30 days for $0.99. This was easy for customers to understand and much easier for call-center personnel to sell. Coordinated marketing and media tactics focused on making it a consumer self-selection. The results were a very efficient increase in vertical service penetration across the entire marketing area.

4. "Get It Together"

Problem: Consumers and businesses do not "shop" call centers regularly for telecom services. They have infrequent contact, primarily for moves, billing issues, and changing services.

Best Practice: **Telecom marketers can greatly increase sales by taking advantage of existing customer behavior, i.e., sell where they are already shopping**. By not trying to change behavior and leveraging existing behavior, telecom marketers can greatly increase exposure and real sales by extending distribution to traditional retail outlets.

Example: To gain top-of-mind awareness with consumers, Sprint formed alliances with retailers that carried products consumers shop for regularly. Radio Shack and mass-merchant electronic stores were utilized to sell Sprint products and services through a "store within a store" merchandising arrangement. Sprint services were offered as a natural compliment to consumer hardware purchases. While this tactic has now been widely emulated, Sprint was able to significantly grow its retail service at the expense of its major competitors. Inspired imitators included Bell Atlantic with Staples and Ameritech with Sears.

5. "Pour Old Wine into New Bottles"

Problem: For telecoms, developing and launching a new product can be a slow process with uncertain results. There are complex system-integration challenges—i.e., ordering, provisioning, billing, etc.—and lengthy regulatory and legal approvals required. In most cases, telecom marketers are selling what they have, not what customers need.

Best Practice: **Bundling/selling existing services in new packages helps telecoms to overcome new product-development challenges and provides a never-ending array of "new ideas."** The trick is how and what to bundle, and how to price it. Customer knowledge and research can provide the answers.

Example: Bell South was an early leader in bundle penetration to both consumers and businesses. They revised their call-center capabilities to emphasize "Complete Choice" bundles versus the traditional à la carte service approach. Another example is Sprint's B2B Long Distance Division. Their "business solution" is a simple package that makes it easy for small businesses to one-stop shop. Both companies have been successful at rebundling existing products to make them a better value, more convenient, and easier to acquire and change.

6. "Develop Firsts That Last"

Problem: Telecom product development is technology-oriented, not customer-focused. As a result, customer acceptance and adoption of new telecom products is slow. Products are created with features in mind instead of benefits. The "build it and they will buy it" mentality must change to improve effectiveness and efficiency.

Best Practice: **In the long term, successful new products and services must be based on addressing and anticipating real customer need, not just leveraging technology.** Again, this requires increased levels of customer knowledge and extensive research.

Example: AT&T developed "Family Packages" that addressed the different needs of multiple consumers in a single household. This increased penetration and revenue per month. Ameritech's "Call Privacy" service was a direct response to a high visibility, controversial concern with aggressive telemarketing to households. Sprint PCS anticipated the issue with wireless signal clarity when they created the "Digital Difference" through their state of the art digital system.

7. "Customer Relationships, Custom Made"

Problem: Telecommunications companies are not known for communicating and building relationships with their customers. High churn rates are based on the lack of, or inability to leverage, individualized customer knowledge. This makes customers feel like numbers, not individuals. Bureaucracy and fragmented company infrastructures contribute to this problem. Telecoms are also poor at post-sell communications. They have minimal dialogue with customers and offer almost no personalization in verbal and indirect communication with customers.

Best Practice: **Successful customer relationships are based on genuine content, true dialogue, and relevance to individual customer needs, not just on "customer appreciation" messaging.**

Example: Ameritech Small Business Services wanted to leverage relationships for upsell and retention purposes. They learned that their small-business customers were very interested and also concerned with technology. Ameritech leveraged this need by creating a magazine/catalog called *Small Business Solutions* to provide added value and information to their existing customers. They added value by customizing the magazine/editorial portion to small-business priorities such as how to get more out of technology and hardware, how to save money on technology and telecommunication needs, and how to finance these needs with flexible payment options through their telecommunications provider. This program generated increased levels of loyalty and added a new profit center vis à vis hardware/software financing.

The Benefits of Best Practices: Ongoing Improvement

While these best practices seem rather obvious, they are not widely practiced in the telecommunications industry. As we have shown, certain companies do certain things well. The key to improving marketing effectiveness and efficiency is utilizing these best practices in everyday marketing activity. Like any improvement endeavor, you can't be good at everything. We recommend that telecommunications marketers pick the ones that are right for their business situation and begin the process by scrutinizing all activities against them.

Getting Better Business Returns from Your MPLS Deployments

Adding Intelligence at the Customer Premises and Service-Provider Edge Is the Key to Making MPLS Efficient and Profitable

Manickam Sridhar

Chief Technology Officer

Sitara Networks

Good Technology, Bad Hype

One article says that multiprotocol label switching (MPLS) technology gives service providers increased control over their networks (*NetworkWorld Fusion*, 7/30/01). The next week, an article in the same publication says that experts are "raising red flags" about the technology (*NetworkWorld Fusion*, 8/6/01).

When any new technology with great promise comes along it is usually accompanied by so much hype that it becomes difficult for service providers to sort out truth from fiction. In truth, MPLS is an invaluable tool in the era of Internet protocol (IP) network convergence, but not if vendors oversell what the technology was designed to do.

To determine whether or not MPLS can contribute to the profitability and efficiency of existing network operations, there are two things to consider. First, what is the service provider's most pressing business requirement? Second, what was MPLS designed and optimized to do?

The answer to the first question is simple: increase revenue! Service providers are under tremendous pressure to move beyond commoditized Internet connectivity. On the surface, this should not be much of a challenge because business customers are clamoring for value-added services. This includes everything from managed services to IP–based virtual private networks (VPNs). The VPN market alone is projected to reach more than $35 billion by 2004 (Source: Infonetics).

MPLS, while it is an excellent technology for improving speed and efficiency in IP networks, is not—*by itself*—an enabling technology for offering value-added services. Here's what MPLS was designed and optimized to do:

MPLS is an Internet Engineering Task Force (IETF) standard designed to speed up IP traffic flow through a routed network. It is essentially a tunneling protocol that emulates an asynchronous transfer mode (ATM) permanent virtual circuit (PVC). Without MPLS, each router in the network has to look at the header of each packet and make routing decisions on a hop-by-hop basis for forwarding the packet. MPLS speeds up this process as follows: As packets enter the network, they are assigned labels that associate each packet to a forwarding equivalence class (FEC), as shown in *Figure 1*. All packets that belong to a particular FEC are mapped to the appropriate label-switched path (LSP) and typically follow the same path.

All subsequent routers use the label as an index to a table that specifies the next hop. The router may replace the old label with a new label identifying the next hop, and the packet is forwarded along. Because there is no further analysis of the packet's header after the initial labeling, the forwarding is much faster and more efficient. Routers equipped with MPLS create local switched paths by simply "swapping" labels throughout the network.

There are several options for marking packets, but the best tool (and the one most complementary to MPLS) is another IETF standard called differentiated services (DiffServ). The network equipment responsible for traffic classification and grooming would use DiffServ code points (DSCPs) to mark the packets so that the MPLS–equipped routers can look at the DSCP in the IP header and map the packet to the correct FEC.

In addition to speed, MPLS provides some of the traffic-engineering capabilities associated with the ATM world. For example, FECs can be established prior to data transmission or can be set up at flow time, which allows service providers to determine routes based on traffic flows rather than the shortest route possible (the way that IP routing protocols normally work).

MPLS labeling is also valuable because it creates a more consistent and uniform way of identifying packets across the

FIGURE 1

With MPLS, an IP Core Emulates an ATM Core—Efficient, Fast, and Manageable

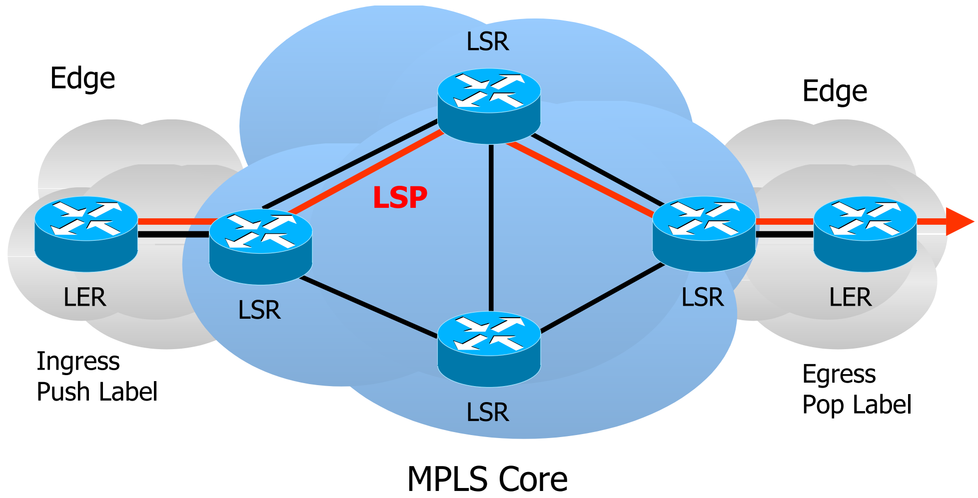

network. Lastly, MPLS makes it easier to deploy larger-scale networks because the label-distribution algorithms learn new paths dynamically.

In fact, the real shortcoming of MPLS, given the service provider's business requirements, is that MPLS alone cannot create premium packages based on a customer's individual traffic requirements.

The Bottom Line Is Granularity

MPLS is designed and optimized to deal with aggregated traffic, rather than individual traffic flows. To maintain the forwarding speed it achieves, MPLS aggregates traffic into broad service categories, or Olympic services, such as Gold, Silver, and Bronze. Traffic can be assigned to one of these service classes and receive the priority associated with each class. However, there is no differentiation possible for individual users or flows once traffic has been aggregated into one of these broad categories.

This means that service providers lack the ability to do the following:

- Differentiate customers within an aggregate class
- Prioritize between a customer's high- and low-priority applications
- Guarantee bandwidth and quality of service (QoS) end to end for individual users and/or applications

Because service providers cannot get this level of granularity with MPLS alone, they cannot reliably offer premium services such as voice over IP (VoIP) and videoconferencing. They cannot guarantee that a customer's mission-critical applications, such as enterprise resource planning (ERP), will take precedence over e-mail or streaming audio. And they cannot guarantee that a burst of traffic from a customer's own users—or another customer sharing the same pipe—will not affect mission-critical traffic.

Granularity Has No Business in the Core

Can the routers at the core cope with the extra burden of more granular classification and provisioning? *Figure 2* illustrates how quickly granular classification and provisioning can cause a nightmarish processing burden at the core of the network as routers try to process millions of flows resulting in *billions* of operations associated with classifying individual flows. This much processing overhead would certainly decrease the throughput of a core router.

The First Touch

To create premium services in an MPLS routed network, traffic needs to be classified and groomed *before* it reaches the core. This "first touch" classification and grooming does the following:

- Enables the core routers to maintain top forwarding speeds
- Provides the deep, granular classification that is needed to deliver value-added service packages based on each individual customer's needs
- Provides end-to-end QoS, which is necessary to offer customers bandwidth and service guarantees

This "first touch" is provided by specialized QoS appliance devices: QoSWorks at the customer premise and QoSArray and QoSDirector at the service provider's data center, as shown in *Figure 3*. This end-to-end QoS solution provides the following:

- Deep packet classification
- Grooming and managing all types of traffic so that service-level agreements (SLAs) can be constructed and traffic volume entering the core is predictable
- Accurate traffic statistics and usage for billing
- Packet marking with DSCP for mapping the traffic into the correct FEC

FIGURE 2

The Job of Doing Fine Grain Traffic Classification and Management Becomes More Computationally Intensive As Millions of Traffic Flows Aggregate at the Core of the Network. The Solution Is to Move This Task of Granular Classification and Management to the Edge of the Network, Where the Traffic Originates.

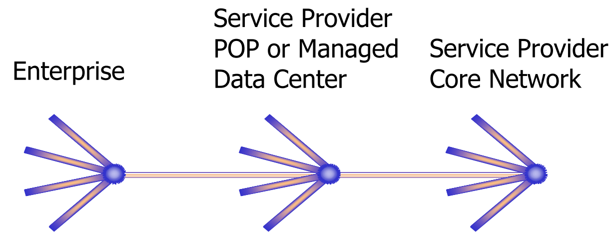

Location

Location			
Aggregation Point	1	2	3
Speed	64K – 100Mbps	10M – OC–3	OC–12 – OC–192
Number of Flows	100,000	10,000,000	1,000,000,000

- Policy management to distribute and track thousands of policies

Deep Packet Classification

By using a combination of the identifiers, shown in *Table 1*, at the edge of the network, QoSArray and QoSWorks provides the deep classification that can bog down edge routers and decrease their performance. QoSArray and QoSWorks automatically classify all IP and non–IP traffic in real time at wire speeds at Layers 2–7. The QoS devices can classify based on default and user-defined parameters and even distinguish between applications using the same protocol.

Traffic Management/Grooming

Core routers equipped with MPLS incorporate basic traffic-management capabilities; specifically, they offer a limited subset of queuing techniques (e.g., weighted fair queuing [WFQ] with weighted random early discard [WRED]). QoSWorks and QoSArray groom all transmission control protocol (TCP) and user datagram protocol (UDP) traffic using a far more sophisticated and complete set of traffic-management mechanisms than routers provide, as shown *Table 2*.

Each QoSWorks at the customer's network is responsible for traffic classification and marking of the DSCPs. In *Figure 4*, all videoconferencing traffic will be marked as higher priority over all other traffic. Videoconferencing packets are assigned to the appropriate MPLS path to deliver true end-to-end QoS.

Service providers are charging an additional 25 percent to 50 percent fee for their premium packages compared with the best-effort offering. A typical service provider could recoup its investment in less than 12 months while enjoying additional benefits:

- Improve customer satisfaction
- Control monthly operating costs
- More accurate billing

FIGURE 3

By Deploying Specialized QoS Devices at the Edge of the Network, Service Providers Can Deliver Precise End-to-End Service Guarantees for Each Customer and Each Customer's Applications, down to the Individual Flow.

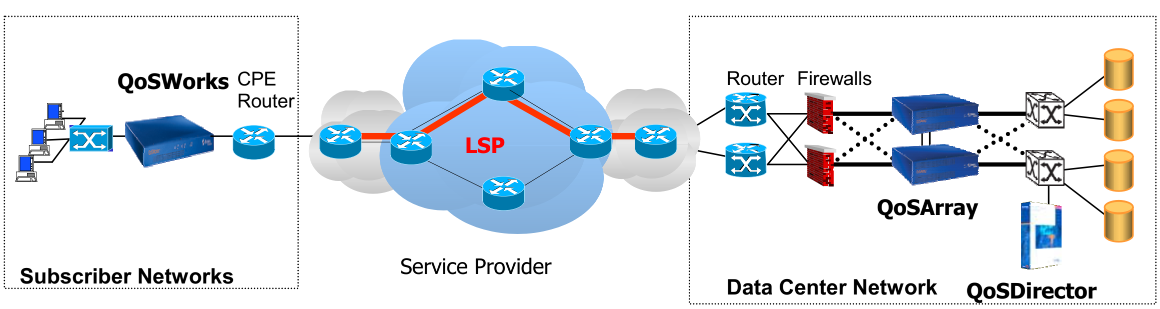

TABLE 1

MPLS Edge Routers Perform Gross Classification by the Following:	QoSWorks and QoSArray Perform Granular Classification by the Following:
Source IP address	Source IP address
Destination IP address	Destination IP address
DSCPs	Source port
	Destination port
	Transport protocol number
	Ethernet protocol type
	Layer-7 state information
	DSCPs

TABLE 2

QoS Mechanism	What It Does	Types of Traffic
Class-based queuing	Provides fine granularity of bandwidth sharing and traffic-priority control, including session level. Enables service-level guarantees for individual flows and aggregate traffic	Works equally well across all TCP and non—TCP protocols. Includes UDP, VoIP, IP, IPX, Appletalk, and SNA
TCP rate shaping	Improves wide-area network (WAN) link efficiency by reducing retransmissions. Provides end-to-end flow control, minimizes queuing delay; and controls bursts	TCP
Fair allocation of bandwidth by connection	Guarantees that all sessions in a policy have equal bandwidth. Important for networks incorporating large numbers of users of session-based applications, e.g., TN3270, CITRIX, etc.	TCP
Packet-size optimization	Reduces packet sizes at the traffic source to any specified size from a maximum of 1,500 bytes to a minimum of 64 bytes. Critical for mixing videoconferencing or VoIP with traffic characterized by larger packets. Minimizes the time packets wait in queues for transmission	TCP
Control bandwidth bursting	Allows traffic classes to borrow idle bandwidth from other specific classes, thus maximizing bandwidth utilization while guaranteeing that applications only use authorized resources	Works for all traffic classes
VoIP controls	Able to identify voice packets from other UDP streams, therefore providing higher voice quality by applying voice-specific QoS capabilities. Includes queue depth and call admission control (minimizes network jitter and packet delays), and prioritization (minimizes packet loss)	VoIP

TABLE 3

	Package 1 Protect CITRIX	Package 2 Deploy VoIP	Package 3 Best Effort
Committed Bandwidth	1 Mbps	750 kbps	384 kbps
Burst Bandwidth	5 Mbps	5 Mbps	None
Customer Priority	High	Medium	Low
Application 1	CITRIX = High Bandwidth = 512 kbps	VoIP = High Bandwidth = 384 kbps	WEB = High
Application 2	WEB = Medium	WEB = Medium	FTP= Medium
Application 3	FTP = Low	FTP = Low	E-MAIL = Medium

A wide range of SLAs can be created to fit a customer s unique requirements. For example, package 1 delivers 1 Mbps of committed bandwidth with the ability to burst to 5 Mbps. CITRIX has a guaranteed bandwidth of 512 kbps and high priority. With Package 3, the customer gets 384 kKbps of bandwidth with no bursting. They are a low-priority customer but can assign a high priority to Web traffic over their other applications.

Policy-Making and Reporting with QoSDirector

Setting policies and reporting on SLA conformance are essential tools for the service provider. Using QoSDirector at the service provider's data center makes it easy and efficient for service providers to create, distribute, and manage even the most complex policies (e.g., by application, time of day, etc.) for hundreds or thousands of customers and tens of thousands of flows. QoSDirector uses hierarchical policy setting, which makes it possible to apply controls at multiple levels of granularity. For example, service providers can structure polices by links, users, applications, or in any combination. Once the policies are applied, they can be monitored and fine-tuned based on real-time performance analysis from QoSDirector. The statistics collected can then be used for checking conformance to SLAs and generating accurate bills.

FIGURE 4

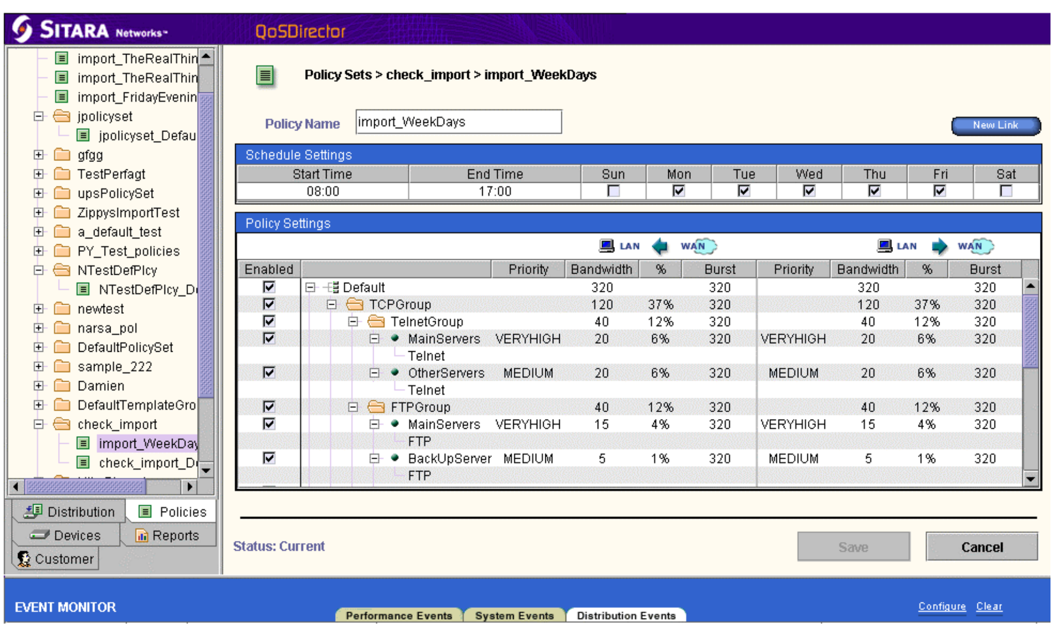

QoSDirector is a standards-based system that interoperates with other network-management systems and the service provider's existing operational support and billing systems.

Making MPLS Pay

When all the hype is over, everyone will get down to the serious business of evaluating MPLS and declaring it a success or failure. If service providers expect it to deliver revenue-generating services by itself, they will be disappointed, but if MPLS does what it does best—making an IP core as efficient and manageable as an ATM core—and service providers add QoS to the edge for classification and traffic management, then tomorrow's headlines will read: "MPLS and QoS Together Provide a Successful Formula for Service Providers to Increase Their Top-Line Revenue."

Appendix: Anatomy of an MPLS Network

A packet arrives at the ingress point of presence (POP) with a destination address marked 192.12.10.6. This packet would typically arrive at the Aggregator at the ingress POP through an access network, and the Aggregator then would forward this packet to the ingress label edge router (LER). The LER looks at a combination destination IP address, the incoming port, and IP packet markings, and then assigns a label to that packet.

This label assignment is based on the negotiations that happened previously between the ingress LER and the core label-switched routers (LSRs). These negotiations can be carried out using one of several techniques, including using the label distribution protocol (LDP), border gateway protocol (BGP) extensions, or resource reservation protocol (RSVP). The end result is a unique label that identifies each packet to a FEC.

For example, Label 5 identifies the unique port that connects this router with the core router. Label 5 is appended to this packet and sent to the core router. The core router ignores the packet's network layer header and simply forwards the packet using the label-swapping algorithm.

When a labeled packet arrives at the core router, the router uses the input port number and the incoming label to perform a match to an outgoing label that is stored in its forwarding table. When a match is found, it retrieves the outgoing label—e.g., Label 7—and replaces Label 5 with Label 7 and sends the packet to the appropriate outgoing interface that connects this router to the next hop in the LSP, which in this case is the egress router.

When the labeled packet arrives at the egress router, the forwarding component searches its forwarding table. If the next hop is not a LSR, the egress router pops the label and forwards the packet using conventional IP routing.

The Transformational Power of Future Wireless Technology

Carol Stephenson
President and Chief Executive Officer
Lucent Technologies Canada

Introduction

This paper provides my perspective on what the future holds in store. First, it will describe some of the major distinguishing characteristics of mobile communications and networking technology in the future. Second, it will take a look at the benefits and opportunities that this new technological frontier presents. And third, it will examine some of the potential dangers.

My message is this: we haven't seen anything yet. With the advent of broadband mobile communications, technology is becoming more personal, more portable, and more pervasive. It is now potentially within the reach—both physically and economically—of more people than ever before. This is where the real promise and the real power of future technologies lie. And it is up to us to realize that promise and leverage its power.

Personal, Portable, and Pervasive

What is meant by the phrase "personalization of technologies?" Personal communications is a term that has been around for a long time, especially in the mobile communications industry. However, we're not talking about PCS or simply putting a digital cell phone in every person's pocket. We're talking about developing wireless communications services that fit individuals like well-tailored clothes.

Increasingly, customers are asking for services that know them, that know what they want, and that know what device they are using. Customers want to be alerted if a danger exists or a new opportunity arises. They want the network to find them or to tell them where they are. Today's new services, such as location-based services, are beginning to meet these needs. But they pale in comparison to what's coming up in the future.

Some of the experts at Lucent's Bell Labs, for example, believe that each of us will soon have our own "cyber clones." These virtual agents will constantly anticipate our information needs, wants, and preferences. They will screen our messages and surf the Web, filtering out irrelevant data and presenting us with the information that we want in the format with which we are most comfortable.

We will also soon be able to customize the way in which we interface with the network in a highly personalized manner at any given time. Software advances—such as "servlets" or executable applications that will run on network servers—will allow us to tailor our interface to the network.

For instance, let's say you're in the car. You can then use your car's technology to interface with the network, or you can use your mobile phone or your personal digital assistant (PDA). At home, you could use your laptop, your gaming device, or your television. The choice would be entirely up to you, because the network would recognize you—no matter where you are or what device you are using.

Generally, network communication will also be more personal because it will be almost like talking with another person. You'll be able to ask the network questions and give it verbal commands. You could surf the Net with spoken words instead of clicking fingers. Voice portals are just beginning to proliferate, but many experts believe that voice will probably be the dominant form of on-line communications in the future. Voice recognition is already available. Voice commands can be used with cell phones to place calls.

In addition to the advent of highly personal communications, advances in microelectronics are shrinking technology—making communications devices, even networks and especially wireless networks, more portable. This technology, known as micro-electromechanical systems (MEMSs) or as systems on a chip (SoC) will likely revolutionize how we use technology. Already, scientists have developed microphones no thicker than a single strand of human hair, engines so small that thousands could fit on a Frisbee, and medical probes that can be threaded through veins.

The MEMS exchange—an umbrella group of MEMS manufacturers—projects that these micro machines will soon generate annual sales of nearly $100 billion in new industrial and medical products. With MEMS technology, scientists believe that we could build satellites the size of baseballs, yet with all the power of giant satellites. And these satellites could be launched with high-powered guns, instead of rockets or the space shuttle—saving hundreds of millions of dollars.

Back on the ground, experts at Bell Labs predict that wireless towers will eventually be so small that they will become part of the optical fiber that will be laid everywhere, bringing truly uninterrupted high-speed network coverage. MEMS will also dramatically shrink the size of communications devices. For example, Professor Clark Nguyen of the University of Michigan believes that MEMS could replace many of the traditional components in cell phones. As a result, mobile phones would become the size of wristwatches, require little in the way of battery power, and would possibly be cheaper. Combine these "Dick Tracy" tools with the new high-resolution monitors that you can clip on to your eyeglasses.

Jef Raskin, author and father of the original Macintosh computer project, tried one of these monitors recently and raved about it. It didn't look like the quarter inch display that it actually was, but rather appeared to be a full-size monitor. And he could sit back, lie down, and move from place to place without missing a Web site. How's that for portable?

MEMS technology promises to make our pockets and purses lighter in other ways as well. In just two years, for instance, watch out for cell phones with built-in bar-code scanners that will allow you to just swipe and buy.

In addition to becoming more personal and more portable, communications technology is also becoming infinitely more pervasive. It is being integrated into machines of all kinds. Bluetooth technology, Wi-Fi networks, and advanced sensors that interconnect machines are just some of the technologies enabling an increasing amount of machine-to-machine and object-to-object communications.

By 2010, Arun Netravali, president of Bell Labs and Lucent Tachnologies' chief technology officer, predicts that this "info-chatter" will surpass the communications volume generated by human beings. And he's not alone. Other experts speculate that by 2004 to 2005, as much as 60 percent of Internet communications will be devoid of human access.

Bell Labs experts at Lucent further predict that by 2025, the entire world will be encased in a communications skin. This skin, fed by a constant stream of information, will grow larger and more useful. This skin will consist of millions of electronic devices—from thermostats and pressure gauges to cameras and microphones—a myriad of SoCs. All of these will transmit data directly into the network, just as our skin transmits a constant stream of sensory data to our brains.

Such systems might be used for anything—constantly monitoring the traffic on a local road, the water level in a river to warn of floods, or even a person's vital signs. For example, a machine in a hospital or at a doctor's office could continually monitor a patient with a heart condition. When something is amiss, that device would automatically alert the doctor. A "pharmacy on a chip" is another possibility. Embedded under a patient's skin, this chip would allow medical personnel to automatically release medication into a patient's bloodstream—even if the patient is at home and not at the hospital.

The Good Things to Come

Personal, portable, and pervasive: these are the characteristics of wireless broadband Internet communications in the future. And at Lucent, it is believed that this new communications capability will rock our world, change our lives, and enable us to do and experience things that we never thought possible.

As Kenan Sahin of Bell Labs explains, "The Internet will evolve from being a complexity that we must spend time mastering to a behind-the-scenes tool that will improve the quality of our lives."

In a recent edition of Forbes ASAP magazine, editor Patrick Dillon is equally enthusiastic from a purely economic point of view. Speculating on the convergence of technologies now taking place, he writes that "looking over the horizon to 2004–2005, we think there really will be a new economy, and this time, we're going to get it right."

According to Forbes, all the signs of a massive turnaround are in place, including the bottoming out of the current high-tech business cycle and the return of money, new technology, and entrepreneurial fervour to markets. In short, the technological revolution will be rebooted. And if history holds true, Forbes predicts that the new mobile and broadband Internet could become a $20 trillion industry by 2020.

It would be nice to believe that some of this prosperity would be shared with the less developed countries of the world—for two very good reasons. First, the new satellite systems and small wireless tower technologies will enable us to build networks that are more cost-effective than any other previous networks. Furthermore, these networks will be easier to build—especially in the remote, geographically unfriendly places that land-based networks cannot reach. Second, the new mobile devices being developed today will not only be less expensive than current devices, they will be easier to use. The Internet will become a visually rich as opposed to text-laden medium. It will become intuitive and voice-based. As a result, mobile high-speed communications could bring the wonders of the Internet to a whole new group of people. It could significantly bridge the digital divide. It could give everyone a voice—a voice of unquestionable value.

Consider what happened in the Philippines last January. Hundreds of thousands of protesters massed in central Manila to oust Philippine president Joseph Estrada. Yet, they were not lured out of their homes and offices by megaphones or gunfire. Rather, they were enticed by millions of instant messages broadcast to their cell phones. Indeed, an estimated 100 million text messages clog the wireless networks in the Philippines every day—making the people of this island country the world's most avid fans of instant messaging.

Closer to home, consider the horrific events that occurred in New York, Washington, and Pennsylvania on September 11. Mobile communications and the Internet were crucial to all kinds of outreach and coping activities—from the cell phone dispatches of survivors to loved ones to the millions of dollars in relief assistance raised through on-line fundraising drives. Some people in or near lower Manhattan were surprised when they turned on their cell phones and were almost instantly contacted by their wireless carrier. The companies were monitoring cellular signals near the collapsed World Trade Center towers to try to find survivors who might have been trapped or unaccounted for. The ability to communicate in the face of such disastrous circumstances was invaluable to countless families, friends, and business colleagues. Their stories now linger in our collective experience. They have forever changed the value that we place on our communications tools.

The Dangers That Lie Ahead

However, the tremendous value of communications technology does not come without a cost. For example, by virtue of their very functionality, broadband mobile technologies will be far more effective in collecting personal information than any other previous technology. How will we strike the right balance between providing personalized service and not intruding on someone's personal business? How will companies protect the privacy of their customers? How will governments safeguard information about their citizens, especially in the face of increasingly porous borders? The new technology will also have the ability to find us wherever we are—perhaps even when we don't want to be found. The distinctions between work and home will further blur. For most of us, this will not be good. We will need to adopt a habit of finding the "off button" and using it.

Indeed, some observers argue that this could prove disastrous. In her new book called *White Collar Sweatshop*, Jill Andresky argues that portable computers, cell phones, pagers, and PDAs have served primarily to extend the workday, rather than radically reshape the nature of work. She further believes that people are working 10-plus hours a day, not because they want to, but because they have no choice. Many companies expect that kind of commitment. And despite that extra work and the fact that more and more families have become dual wage earners, many people are not gaining ground.

According to the Ottawa-based Vanier Institute of the Family, for example, the average Canadian family has not seen an improvement in their living standard for about 15 years. How can we ensure that broadband mobile Internet communications makes people's lives better? How do we make sure that it enriches work—that it makes work more rewarding and more creative—not just more productive? These are the hard questions that we must find answers to. If we do not, the broadband mobile Internet will never achieve its true promise.

Conclusion

To sum up, we are about to enter a new technological age—an age of personal, portable, and pervasive technology—the age of broadband, mobile Internet communications. This age will radically transform our lives. It will bring us services customized to our individual needs. It will make technology easy to use anytime, anywhere. It will allow technology to anticipate our every need and to truly make our lives easier.

This next wave of technology could kick-start the economy. It promises to bring everything the new economy was supposed to bring in the first place: new industries, profitable business models, and enduring innovation. Potentially, it could give everyone on this planet his or her own voice for the first time.

But, there are dangers inherent in this new age. The new technological order could jeopardize our privacy, erode our individual senses of identity, and create unforeseen personal and social pressures for people—both on and off the job. Like all great technological innovations of the past, the era of broadband Internet communications will create tremendous opportunities and equally daunting challenges. And, in the end, its success or its failure will come down to the decisions that we make today.

The biggest and most difficult of these decisions may concern the type of legacy that we would like to leave behind. Throughout history, the great business people have always done more with their good fortune than simply build a business and make a profit. During the mid-1500s, for example, Cosimo de' Medici of Florence, Italy, built and nurtured the greatest banking dynasty that the world has ever seen. Yet at the same time, he cultivated an enduring legacy of patronage to the arts—a legacy that ignited the Renaissance.

Andrew Carnegie made his great fortune in steel, railroads, and steamships and, in 1901, pocketed $240 million when J.P. Morgan bought him out. But Carnegie then set out to distribute his wealth, most notably by providing public libraries throughout the United States and Great Britain.

Henry Ford's dream was to make an automobile that everyone could afford and in the process of achieving his vision, he became history's first billionaire. He also paid his workers more than any other American manufacturer. He introduced the eight-hour workday and instituted a profit-sharing plan that would distribute up to $30 million annually among his employees.

How will we, as business leaders, make sure that the world we create through broadband, mobile communications will be a better one? How will we work to ensure that this new technology not only improves our bottom lines, but also brings the wonder of Internet communications to people who might never have the opportunity without our help? How will we make sure that the broadband, mobile Internet improves the quality of life for our employees and their families?

These are some of the difficult questions that we should ponder. These questions are difficult, but we must ask them. Our answers to these questions today will be our legacies tomorrow.

Standardization in Fixed Broadband Wireless Access

Paul Struhsaker
Chief Technology Officer
Raze Technologies Inc.

Since 1995, consumers' demand for high-speed Internet access and data-intensive broadband communications has grown exponentially, both at home and in the work place. Clearly, we are still in the early stages of deployment of broadband access. There are more than 100 million households in the United States, and, by the end of 2001, the rollout of both cable modems and digital subscriber line (DSL) will have provided access to less than 10% of the market. The Federal Communications Commission (FCC) in a 134-page study predicts that by 2004 nearly half of all Internet connections in the United States will be high-speed broadband.

Until last year, it had been expected that cable modems and DSL could provide service to as much as 90% of these households. That projection is being significantly reduced due to impairments and limitations in both the cable and voice physical line plant and due to accompanied deployment issues that have surfaced in the last year. We increasingly hear about a growing "digital divide" that exists between larger metropolitan areas that are seeing deployment and infrastructure upgrades and smaller cites, towns, and rural areas. The digital divide is significant and affects from 35% to 50% of the available market.

Fixed wireless broadband access and rural service using satellite broadband services are increasingly seen as the solution to address these markets. With the continued spectrum reallocation efforts by the FCC, nearly 300 megahertz (MHz) of licensed spectrum and 100 MHz of unlicensed spectrum will be available by 2002 to 2003. Cost-effective broadband wireless has the ability to be a viable solution to address this underserved market.

The key issue for the fledgling fixed wireless industry is the development of interoperable standards so that fixed wireless can achieve the economies of scale that DSL and cable modems have achieved through the respective standardization efforts of their industries.

It is beyond the scope of this paper to address all standards efforts in broadband access. The remainder of this article will examine both the benefits of standards and the efforts under way to achieve interoperable fixed wireless standards and concerns that stem from these issues.

Standards allow equipment from multiple vendors to interwork seamlessly. They also promote competition and, in general, reduce the operating costs for all. Once a standard is established, everyone in the supply chain can streamline production and reduce expenses. The chip/component manufacturers can reduce costs by mass-producing their products. Equipment manufacturers can reduce expenses and product prices by purchasing from multiple sources, knowing that each component is built to a standard. Similarly, carriers reduce their costs by integrating equipment from competitive, standards-based products that have predictable interfaces and functionality. This chain of cost-efficiency results in reduced prices, higher-quality products, and significantly increased competition and choice for the consumer.

Some classic examples of this can be seen in the wireless industry. When cellular wireless was first introduced, there were a few proprietary technologies. The early deployments of these technologies resulted in disparate networks, inefficient equipment, and high consumer prices. As the industry developed common standards (e.g., code division multiple access [CDMA] and the Global System for Mobile Communications [GSM]), the prices dropped and the use of cellular phones increased to the point where today cellular phones are almost ubiquitous. Another example is wireless local-area networks (WLANs). In 1993, WLAN was a limited market due to the high cost of non-standard, proprietary equipment with products costing well over $500 and providing less than 128 kilobits per second (Kbps) data throughput. In collaboration with the equipment suppliers and silicon vendors, the Institute of Electronic and Electrical Engineers (IEEE) established the 802.11 standard for WLAN technology. Today, standardized WLAN products have been reduced to ~$75 and deliver ~11 Mbps data throughput. Enhancements to 802.11 will result in ~50 Mbps products reaching the market in 2002.

The penetration for broadband access to the residential and small office/home office (SOHO) is highly price sensitive. There are more than 70 million Internet users in the United States. Of these 70 million users, more than 52 million are dial-up users whose access charges are on average $25 per month. Several studies of the cost sensitivity of consumer interest in broadband Internet service directly relate penetration to pricing. As *Table 1* indicates, the price must fall to $25 per month (the same price that people currently pay for dial-up service) for broad market penetration.

Without standardization and the associated cost-reduction factors that will result, service providers will not be able to meet consumers' pricing expectations.

Clearly, service providers in fixed broadband wireless access (FBWA) want three basic issues solved by their equipment vendors: subscriber cost, radio frequency (RF) link performances (e.g., the non-line of site problem), and standards to allow the use of multiple vendors. The current generation of proprietary FBWA solutions fail at all three of these critical measures.

If standards are so important and are a requirement to reduce cost, why are the fixed broadband wireless standards efforts so fragmented?

One aspect is economic self-interest. Early market entrants and/or major corporations would like to make a de-facto standard using their technology. The success of Qualcomm in the establishment of CDMA cell-phone technology is the best model. Qualcomm was successful due to the demonstrated superiority of their technology. Where Qualcomm's introduction of CDMA for civilian use was a revolution (even though it had been used for decades in the military), FBWA technology is an evolution. FBWA is based on the synthesis of a number of existing technologies such as optical frequency division multiplexing (OFDM), block frequency equalization, adaptive modulation, block adaptive error correction, adaptive beam forming, and multiple channel processing. All these technologies have been field-proven and are decades old. It is their low-cost integration to create the next generation of FBWA products that is unique.

The second aspect is the wider set of requirements that FBWA systems must address. Spectrum for FBWA that is applicable to low-cost home and SOHO deployment will be in four separate bands between 2.5 gigahertz (GHz) and 6 GHz in the United States. One of these bands can support isolated uplink and downlink channels such as the cellular technology frequency division duplex (FDD). The other bands provide uplink and downlink on a single channel similar to WLAN products (e.g., time division duplex [TDD]). Some service providers need support for data and voice (primary voice in some instances), while others require only data. FBWA must be able to evolve its level of service in the future to provide value-added products that provide incremental revenue to service providers. Toll-quality voice continues to be the single most important value-added revenue for any operator. As data revenues erode, voice, streaming media, and other quality of service (QoS) capabilities will be required for operator and service providers to generate healthy revenue and profits.

The efforts to create a fixed wireless broadband standard can be divided into two types of organizations: open and closed.

Open standards groups are accredited by national and/or international standards organizations. Representation is based on individual attendance and participation by written submission. These organizations use due process and voting to achieve consensus in development of the standard. Group/block voting by a single corporation is prohibited.

TABLE 1

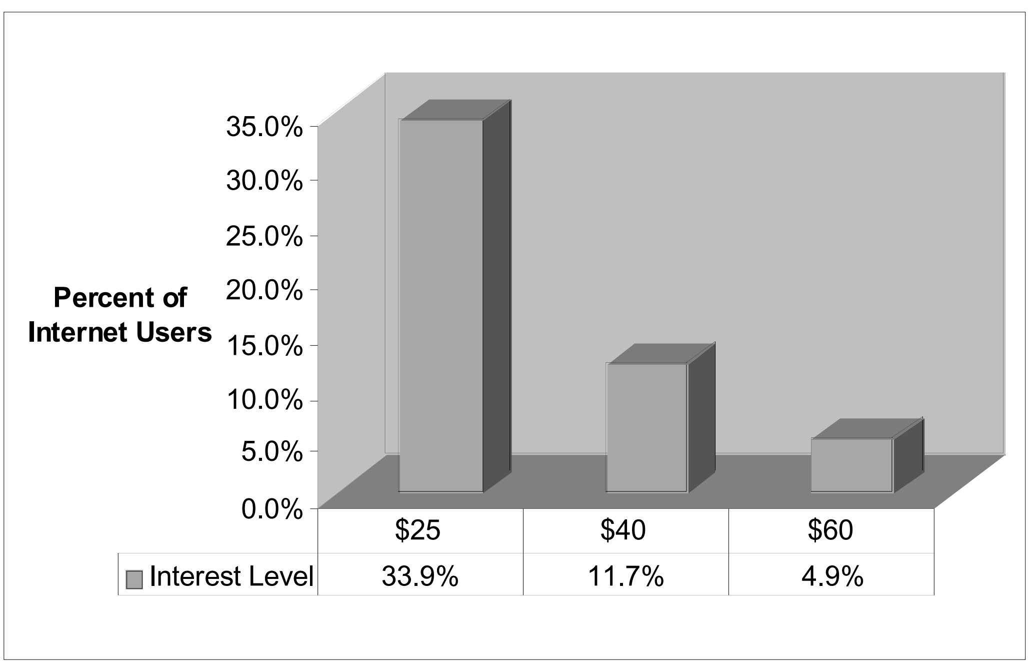

Most important is that the open standards provide a right of appeal inherent in the standards letter ballot process (More information on the IEEE Project 802 is provided in Appendix A.).

The two principle open standards groups in FBWA are as follows:

- IEEE 802.16 Broadband Wireless Metropolitan Networks @ www.ieee802.org/16/, which consists of four task groups:
  - TG1 frequencies greater than 11 GHz
  - TG2 coexistence/RF planning (a function lacking in all other standards bodies)
  - TG3 frequencies from 2 to 11 GHz
  - TG4 unlicensed frequencies (UNII 5.8 GHz band in the United States)
- ETSI BRAN (Broadband Radio Access Networks) @ www.etsi.org

Closed standards forums are corporate based. Voting is at a corporate level and members have unequal voting rights. Furthermore, these groups have executive power that has a concentrated vote and can force the direction of the group. In some cases, assignment of intellectual property rights is a requirement to become a member.

The three principle closed standards groups in FBWA are as follows:

- BWIF (Broadband Wireless Internet Forum) @ www.bwif.org, chartered by Cisco Systems to promote their current FDD–OFDM product as a standard
- WDSL (Wireless DSL) Consortium @ www.wdslconsortium.com, which is more of an industry group that is actively participating in the open IEEE 802.16 standards process
- OFDM Forum @ www.ofdm-forum.com, which is promoting a patented technology for OFDM from Wi-LAN, is also participating in the IEEE 802.16 standards process

There is, today, a great deal of criticism of the progress made in open standards bodies. That criticism is unwarranted. In the latest 802.16 session, the IEEE 802.16 TG1 is completing the ballot stage of the process and will soon (Q4 2001) be an approved standard. At the same meeting, TG3 completed the basic draft standard, and that draft is currently undergoing a formal comment process. Making extensive use of the TG1 media access control (MAC)–layer standard has accelerated the TG3 efforts.

Both the WDSL and OFDM Forum industry groups are active participants in these IEEE 802.16 TG3 efforts. To further accelerate the standards process, IEEE 802.16 is holding six meetings a year, as opposed to the standard three meetings, to increase progress. In addition, e-mail reflectors are being used to reach consensus on technical issues between sessions.

TG3 is expected to have a draft ballot specification in the fourth quarter of 2001 and ballot to begin in the first quarter of 2002. Demands on service providers and actions by the FCC are also helping to drive standards to move faster.

Compared to closed standards, IEEE 802.16 is making a series of technical and technology cost trade-offs to ensure that the long-term cost structure of products will remain low. In particular, the subscriber costs and complexity are being reviewed. 802.16 is also addressing a common MAC and basic physical-layer (PHY) structure that will support all TDD and FDD bands in the 2 GHz to 11 GHz band. Like IEEE 802.11, we can expect a series of standards enhancements from IEEE 802.16 to address this issue.

Which leaves us to reconsider our question: Why do standards activities appear to be fragmented in broadband wireless?

In fact, they are not fragmented. Consensus is being rapidly built on a much wider set of standards issues that address a world market—not just a U.S. domestic market. Consensus will be reached at a more rapid pace, as the standard will leverage high frequency standards efforts that are completing the ballot process.

In spite of the consensus being built, why is BWIF the only forum that is not participating in the open standards process? The answer is simple: Because they want to be considered a de-facto standard. Unfortunately, the current BWIF architecture can address only less than half of the available domestic frequencies that will be available in 2002. Modification of their baseline architecture will result in the same issues internationally.

In the end, the open standards process produces a broader standard that will see continued enhancement and acceptance. The worldwide success of projects 802.3 and 802.11 have proven that, in the end, open standards provide the key to better services, broader support, and continued product improvement.

Appendix A

IEEE 802

The IEEE project 802 local-area network (LAN)/metropolitan-area network (MAN) standards committee (LMSC) held its initial meeting in February 1980. It was sponsored by the IEEE computer society to create a single 1 Mbps to 20 Mbps LAN standard. This and subsequent meetings divided the task of creating a standard into three distinct layers:

- The physical layer (PHY)
- The media access layer (MAC)
- Higher-level interface (HILI)

In practice, project 802 develops standards for the lowest two layers (PHY and MAC) of the open systems interconnection (OSI) reference model. A guiding principle of the group has been to have a common interface at the logical link control (LLC) sub-layer. This ensures that standards efforts at higher layers in the OSI model will have common support by all the supported standards.

During the 21 years of operation, project 802 has provided project authorization (PAR) for a number of working groups and their subordinate technical advisory groups (TAGs). The following is a list of the current active working groups:

- 802.0 Sponsor Executive Committee (SEC)
- 802.1 High Level Interface (HILI)
- 802.3 CSMA/CD for 10/100/1000/10000 based LANs
- 802.11 Wireless LAN (WLAN)
- 802.15 Wireless Personal Area Networks (WPAN) [Bluetooth and derivative]
- 802.16 Broadband Wireless Access (BBWA)
- 802.17 Resilient Packet Ring (just authorized)

IEEE standards are coordinated with both U.S. and international standards groups and a majority of IEEE standards are submitted to, and approved by, the International Standards Organization (ISO). Once a project has been authorized (an approved PAR), the working group goes through a series of technical submissions, discussions, and votes to develop a draft standard. Once the initial draft standard is written, a series of technical and editorial evaluations and changes are conducted. This culminates in a letter ballot where the voting members formally approve or provide specific written amendments to the letter ballots.

After the letter ballot has passed and the amendments have been addressed, the draft standard is provided to the IEEE Standards Board Review Committee. Upon approval of the IEEE, the standard can be published as an official IEEE standard. It is now common practice to have a parallel submission to the ISO JTC1/SC6 (Joint Technical Committee 1, Sub-Committee 6, responsible for LANs).

IEEE 802.16 is seeking to leverage approval of 802.16 TG1 (frequencies greater than 11 GHz of operation), which is entering letter ballot to greatly accelerate TG3 (frequencies from 2 GHz to 11 GHz of operation).

Project 802 holds plenary meetings three times a year. For more information, please refer to the IEEE project 802 Web site at www.ieee802.org.

Next-Generation Services: A Service Developer's Perspective

Arshad Syed
President
Integrated Telecom Solutions

This paper offers an overview, primarily from the service-developer perspective, of next-generation services, user solutions, and convergent solutions. Although the opportunity is tremendous, consideration must be given to cost constraints, especially for companies with strict budget guidelines and for smaller companies that cannot afford a great deal of speculative development. In the United States alone the number of wireless subscribers is estimated at more than 100 million, with an estimate of nearly 120 million by the end of 2001. Enormous growth potential exists in North America alone with even greater potential worldwide. The market demand definitely exists. With regard to the speculative opportunities, the time-to-market drivers must be balanced against the fact that, when a service is developed, probably only a very small segment of the marketplace is available because international ubiquitous services do not exist today, although, hopefully, they will in the future. If a service is developed for the United States, for example, it will be applicable only to North America for the most part and not worldwide. The return on investment (ROI) always must be balanced against this other business driver to the realistic short- and mid-term business case.

Service Drivers

In the past few years many services that were called next-generation services had very disappointing take rates for the carrier, the big vendor, and the smaller companies. Service providers and network operators want the best possible service and the most reliable system, and, of course, they want it at a reasonable cost. Significant pricing pressures exist. Future technology and the next generation, however, hold the promise to make these things possible because the volume will drive all these services and the revenue. Consequently, consideration must be given not only to the estimated 120 million U.S. wireless subscribers but also to the worldwide market, which is the true mass market. One of the technology drivers that will have a direct impact on profits and services is broadband wireless. Some of the advance-enhanced technology, such as wireless handsets with nice visual displays and the bandwidth that will be in the network to support it, is very attractive from a service-provider perspective. Mere proliferation of the Internet and the wireless Web will not improve the take rate or the usage, but services that offer greater value and increased productivity will be a key factor to better penetration.

Service Solutions

Presence and actual location—which will soon be available in North America—coupled with the Internet, some of the bandwidth that will be in the network, and basic wireless intelligent network (WIN)–type standards could lead to valuable solutions. Some will perhaps be a retooling of existing solutions and others will be brand-new solutions. Parlay and some other standard application programming interfaces (API) may allow providers to develop one platform that then can run on different types of infrastructures and networks.

Global network interoperability, of course, has been both positive and negative depending on the perspective taken. Progress has been made on standards, but in today's world there is a duality of networks and a duality of standards. Roughly half the world is Global System for Mobile Communications (GSM). North America and most of South America are American National Standards Institute (ANSI)–41 compliant. The same is true for advanced intelligent network (AIN), capability set (CS)–1, and the International Telecommunications Union (ITU) standards. There is a lack of interworking gateways even though these standards have been out for a long time. No major service provider really has installed gateways that would interwork within, for example, the United States and the United Kingdom, or the United States and Japan, or Brazil and Germany. Significant traffic flows exist among all these countries, and those gateways could be valuable, but most carriers currently are so focused on data that this is being overlooked. Those gateways could allow international ubiquity of service rather than only roaming. The national and international variances do cause problems occasionally. While service solution providers often are expected to do speculative software development, more leverage, especially with some of the new technologies and some of the initiatives, would expand solutions development.

The convergence of network standards, the third generation (3G), and other developments are very promising, as is the set of new standards from the Internet protocol (IP) side. For example, a traditional AIN or WIN, when coupled with IP, generates some very interesting possibilities. Some of the standards that are being adopted quickly on the IP side are worldwide and global. The Java and the hypertext markup language (HTML) or wireless markup language (WML) are the same in the United States or in an Asian country. For instance, the secure sockets layer (SSL), Java, enterprise Java, and common object request broker architecture (CORBA) are all the same, which is encouraging. Standards make it possible to think about solutions and develop them. In general the industry's move toward one worldwide standard is very positive, facilitating innovation, creation, and competition.

Regulations

Regulations, although often perceived as negative, actually can be positive on occasion. For example, the enhanced 911 (E911) mandate issued several years ago will open new location services and enhanced location services because, once location-finding equipment and the infrastructure are deployed, operators will want more mileage beyond E911. Local number portability (LNP) is another example of a regulation leading to innovation and new service deployments. Recently, the Federal Communications Commission (FCC) issued more regulation in wireless caller name and number portability which will likely result in more service deployments and increased value for the consumer. Thus, regulations actually are sometimes very helpful and positive.

Evolution versus Revolution

In terms of the infrastructure challenges and opportunities, for some operators with significant sums invested in infrastructure, it will be difficult to abandon those systems and go to wireless or broadband wireless quickly. Thus, new services must be able to exist within the current infrastructure and evolve with it for a period of time. Some of the capabilities, such as broadband and the Internet, are interworking and eventually will drive mass usage that can result in lower user costs and lower operating costs. There is potential for many millions of wireless users, even billions of wireless users. The existing technologies have to be leveraged, and a strategy must be in place to bridge to the next-generation networks (NGN). There will be a gradual evolution, rather than a revolution, to the NGN and to next-generation services.

With regard to infrastructure, some of the current services are very popular and will remain in demand. Those services will continue to exist for a while. Therefore, new services must coexist and interwork with the old services. What will be needed is a means of taking some of the old solutions and retooling them to use the Internet and some of the more available bandwidth. As a result, global access in a seamless manner definitely will drive the usage.

Next-Generation Architecture

Convergent architecture will need support for new and traditional interfaces as well as support for IP, best-price performance, and ease of use from a subscriber's point of view. In addition, support for these must be built into the infrastructure and platform—neither as an afterthought nor only as the responsibility of the service or the solution that will run on it. There also will be demand for high reliability and high availability along with flexible characteristics based on the applications. Running an application that is mission critical, for example, will demand higher reliability, and a greater price/performance ratio will be acceptable for that. Running an enhanced service that does not need to be available continuously must be able to take the same platform but with a lower reliability and availability, thus a lower price/performance ratio. These flexibilities must be available.

Migration to next-generation architecture should be driven by services rather than by technology. This is a key differentiator. It is possible to implement the technology and deploy it without knowing how it will be used. At the same time, knowing the kind of services desired could interest subscribers and drive costs down while increasing access. The global 3G migration, for instance, is very important. Ubiquitous services should be driven by content and by personalization. Personalization in this sense means knowing the user, the user location, and the user profile and preferences so that this profile can be applied in real time to increase productivity and add value.

Service Evolution

Evolution on the services currently available—caller name, caller identification, virtual private network (VPN) on the wireless side, mobility, and so on—requires that these services be opened to the Internet. All the services that are deployed today will continue for some time, and the life of those services needs to be extended by opening the access to the service data for subscribers using the fixed Internet or the wireless Internet. This is one way of bridging, or moving toward, the next-generation services. Another strategy is retooling of services for better value by leveraging IP technology. IP can be used to retool many services and cut the cost of deploying and operating services for the network operator. Opening service data for everyone to use on the Internet, in a secure and simple manner, and retooling some of the services, would lead the way to fully convergent services.

Future Scenarios

A number of services might be seen as next-generation services, or even convergent services, because they use aspects of telephony, wireless, location, and the Internet. For instance, location can be coupled with incoming call treatment. A user's only number will then be the wireless number. Depending on the user's location, calls are automatically sent there—home, office, or elsewhere—because the presence of the user's handset is recognized there. Calls automatically will go to voice mail if the user is at a place where it is indicated the user does not want the phone to ring. Similarly, when the user is not close to any locations that the service knows about, it will ring the wireless telephone. Location can be used in this way by coupling existing wireless networks and the fixed network to offer a service that is attractive and valuable for subscribers.

Caller identification or caller name also can reach improved value by applying new technologies to it. Instead of a user having to call voice mail to check messages, the termination attempt trigger, which contains all the information that will be displayed in a caller identification box at home or office, can be used to provide a notification on a wireless telephone in real time when a call is made to home or office. In addition, a personalized and secure Internet portal will exist that shows what happened and who called. While checking e-mail on the Internet, for example, a user could see a screen pop up that would indicate someone is calling. If this service was implemented on a softswitch, that call could potentially be taken as voice over IP (VoIP) from a laptop or on a wireless telephone because those capabilities may exist to connect those calls. This portable caller identification also will have the capability to show a call in progress. For instance, if a user misses the screen pop, the call will go through because it is in real time, and perhaps someone will answer the phone. After the fact, because it is the softswitch service, it will show an updated line of text indicating that the call is in progress, and the user could elect to join in on that call as a reverse of three-way calling. This is an example of the possibilities once IP, wireless, location, and other elements are brought together.

Front-End Communications Processors: Their Place in an IP World

Angus Telfer
Chief Technology Officer
INETCO Systems Limited

Introduction

Communications front-end processors (FEPs) are responsible for linking client applications and their associated networks to host computer-based applications. With the advent of the Internet and of Internet protocol (IP) as a universal protocol, it is often assumed that there is no longer any need for FEPs as "everything is IP." This may well be true where FEPs provide only straight connectivity (and assuming IP never changes). However, FEPs also perform a number of other vital communications-related functions that are closely linked to transaction applications, including message and transaction switching, multiplexing, transaction security, quality of service (QoS) guarantors, and end-to-end transaction management and reporting. The need for these functions is especially important in mission-critical transaction environments, such as point-of-sale, security, and health-care applications. In these environments, FEP functionality is more necessary than ever before.

This paper examines the functions performed by FEPs in relation to IP network environments. It also outlines different approaches that may be used to implement this functionality and discusses the advantages and disadvantages of each.

FEP Functionality

FEP functionality is primarily at the communications level. In terms of the open system interconnection (OSI) model, this means the transport, network, protocol, link, and physical layers. Functionality provided can be broken down into the following categories:

- *Connectivity:* Involves the connection of networks, devices, and servers
- *Switching:* Involves the switching of packets and messages based on network addresses and message contents
- *Multiplexing:* Involves the multiplexing and demultiplexing of data streams, such as the aggregation of transactions in retail point-of-sale applications
- *Reliability:* Involves issues such as end-to-end delivery confirmation, maintaining a necessary QoS, load sharing, and appropriately pacing transactions so as to not overload any device or server
- *Security:* Involves data-security issues such as encryption, firewall type issues such as the examination of data address information before routing packets, and device-security issues such as ensuring that there is no way to penetrate and take control of a device remotely
- *Network Management:* Involves network monitoring and network control
- *Reporting:* Involves the reporting of network and communications activities for analysis. These reports may be used for customer billing, network planning, or for proactively finding and eliminating network problems.

Each of these categories, and how it pertains to the IP world, is explained in the following sections.

Connectivity

A traditional use of FEPs is to act as a multiple protocol gateway for connecting diverse equipment. This is how it got the name "FEP" in the first place. It "front ended" the mainframe computer.

At first glance, this function is no longer needed in an IP environment, as all devices appear to be converging to a single communications infrastructure. However, things are not that simple in reality. New IP–based protocols such as wireless application protocol (WAP) and short message service (SMS) are required to support wireless "always-on" applications. Other new IP protocols, such as test transmission control protocol (TTCP) and QTP, are required to properly support transaction-based services. Finally, even existing protocols, such as IP, hypertext transfer protocol (HTTP), etc., are being evolved to meet increased requirements for security, addressing, reliability, priority traffic, etc. The result is that the IP infrastructure is steadily moving to having as many (or more) different protocol implementations than there were in the "bad old days," where each equipment manufacturer supported their own suite of protocols.

Integration of these new IP protocols and protocol variants over time could presumably be done by upgrading all equipment simultaneously. However, this is not practical. This is especially true in these days of network appliances where each appliance performs specific, well-defined functions. Even if it were possible within a particular enterprise, connection to the network outside the enterprise would necessitate a multiple protocol gateway of some sort.

Switching

Another important traditional use of FEPs involves the switching of packets and messages based on network addresses. In modern implementations, this may also be based on message contents. One message may be sent to one server while another message may be sent to another server based on such items as the "terminal identifier" or "bank identifier" field within the message.

With the Internet, switching is accomplished at the client (e.g., PC) end of the network. Control of all communications, including making simultaneous connections to several destinations, is clearly left with the client (e.g., Internet browser, e-mail, etc.). While this approach works quite well in client-centric systems, it is potentially disastrous architecture for mission-critical applications where the enterprise must maintain constant control over the device. With a client-centric architecture, there is no reliable way to manage, monitor, or audit the device in real time, as it may be communicating with several remote servers simultaneously. As well, such an approach provides a large opening for hackers to steal control of the device.

The answer is to have a communications processor (e.g., an FEP) through which all the device's communications goes and through which it can be monitored, audited, and controlled. This device can then route the messages appropriately and can also be interrogated at any time with respect to the status of any device, as well as being able to report any abnormal events that occur.

Multiplexing (Aggregation)

One particular type of FEP that is used within many financial networks is a network access controller (NAC). As well as the other functions described in this paper, a NAC is capable of concentrating multiple data streams onto a single data stream by acting as a gateway between circuit-switched data and message-switched data. This multiplexing (or aggregation) of data allows there to be fewer circuits on the host side of the NAC than there is on the device side. It also means that the host side circuits can be static as opposed to the device side, where, in the transaction world, circuits are often transient and generally last only a few seconds each.

This function is less important in the IP world than it is in the legacy world, where circuit management traditionally placed a large load on the mainframe. However, it is still important in large networks, as transmission control protocol (TCP) connections are still relatively expensive and the maximum number of connections into a server is generally much less than the number of concurrent transactions that could otherwise be handled by that server.

Reliability

Reliability involves issues such as end-to-end delivery confirmation, maintaining a necessary QoS, load sharing and transaction pacing so as to not overload any device or server, and alternate routing in case of server failure.

Traditionally, the FEP would be responsible for these functions and would handle them as follows:

- Message delivery would be acknowledged end-to-end. That is, a message would not be acknowledged back to the application until confirmation was received from the remote device either by a direct acknowledgement or a delivery-confirmation message.
- QoS information would be maintained on an end-to-end (e.g., client device to server application) basis by measuring the time before a confirmation was received, by noting retransmissions and other recovery actions, and, if necessary, by alerting the network operator.
- Load sharing and transaction pacing was done by the FEP holding onto transactions and spreading them over available circuits (e.g., via rotary hunt groups, etc.) so as to not overload any device. Channels to the host could be "paced" by allowing transactions through at only a specified rate. Alternatively, channels could also be configured as half duplex (i.e., one-way only) to prevent the turn-around delays encountered in many hosts.
- Failure recovery was done by the FEP transparently routing transactions to an alternate host should the primary host fail. Note that it is generally important in these scenarios that all transactions be routed to either one host or the other in order to maintain the operational integrity of the application and its database. If the client device controls the switching between hosts, it is invariable that some devices will be connected to the primary server and that some (due to purely transient issues) will end up connected to the secondary server. At this point, the enterprise is, whether it wants to or not, running a distributed database.

In pure Internet-based IP networks, some of these functions are provided by the router network, while many are not done at all. This results in a poor (and generally unknown) QoS. In addition, the host must be thoroughly tested in advance to ensure that it can always take the simultaneous, worst-case transaction hit, as there is nothing in the network to spread load or to pace the rate of incoming transactions.

Security

Security in IP networks is primarily accomplished by the use of firewalls and virtual private networks (VPNs). These approaches provide a level of security but do not eliminate the possibility of insider attacks on the network. For example, a client device could be being taken over by hackers, or an attack could be mounted on the server.

Traditional FEPs do not provide much in this area other than the examination of addressing information before packets are sent on. However, FEP–like functionality is increasingly needed that would allow the client to be blind to all incoming requests except those from the FEP. While this is not strictly in tune with the concept that everything should be simple network management protocol (SNMP) manageable, should respond to pings, should allow downloads by file transfer protocol (FTP), etc., it is much more secure than the approaches used in most current IP networks.

Another instance in which FEP–like functionality is required is the filtering out of messages that do not conform strictly with expected transaction types before they get to the host. Firewalls currently do this to some extent. As mission-critical applications expand on IP networks, it will become increasingly necessary to expand this capability to the actual transaction message itself.

Network Management

One of the areas in which traditional FEPs excel is in their ability to monitor the health of all network elements, including that of the client devices themselves. The FEP is in the middle of every message and so always knows the state of the device without the need to poll it.

In most current Internet-based IP implementations, there is no central network entity. Instead, network-management functions are generally placed off to the side, where an SNMP server such as HP OpenView polls devices, receives alarms, and notifies the operator as it deems appropriate. Putting network management off to the side like this is, at best, a very poor substitute for providing in-line network management. It eliminates the possibility of using network management as part of a dynamic feedback system where corrective action can be taken automatically in real time. It also increases the bandwidth required due to the management requests that must be made over the network.

Reporting

Reporting on communications activities is important, as it generally affects core business operations, such as network planning, billing for network use, auditing (for security), and the interception of data as required by some governments for digital wire tapping.

Centrally located FEPs are an ideal place to collect the data required for such reports with minimal network load or changes to server applications. The alternatives are to collect the information from the devices themselves (impractical for large networks and for places where the devices are owned by third parties) or to collect the information from the server (impractical where multiple server applications are involved or where the servers are owned by third parties).

FEP Implementations

As can be seen, FEP functionality is very important in IP networks used for mission-critical transaction processing. As our dependency on real-time transaction systems grows, it will become increasingly vital. Given this, how can such functionality be achieved? Following are some of the methods that have been used.

Client as FEP

One way that has been promoted is to move the responsibility to the client as is done with Web browsing over the Internet. In this model, the client, such as an automatic teller machine, connects to each server application directly. That is, it connects to one server for financial transactions, another server for device management, another server for network management, another server for providing local advertising, and may even connect to other servers for such things as requesting cash replenishments, etc. In the case of network or server failure, the client simply finds another server to connect to independently of any central control.

The primary advantages of this solution are in favor of the client manufacturers. It is simple to implement, and it moves the control of the device from the central server application to the client. This reduces the reliance of the client manufacturers on the producers of the central servers and allows them to make their client application more proprietary.

Key disadvantages of this approach are as follows:

- The client is operating essentially independently of any central application. This makes it very difficult to control or monitor. If security breaches do occur, they can be difficult to detect, and it can be very difficult to regain control in order to limit the damage or correct the problem. This lack of a central control or audit point can be disastrous in mission-critical environments.
- As the clients operate independently, there is no good way to switch them from one server to another server in any synchronized manner. In normal operation, clients will tend toward being distributed over all available servers due to variable loading conditions on the network. This is not a problem if servers are identical and meant to be operated in this fashion. However, in the case of most enterprise servers, there is the concept of primary and secondary servers (if for no other reason, then to provide a method of updating server software). Distribution over these servers without appropriate operator control can result in problems maintaining consistency within the database, with encryption keys, etc.
- As all clients are operating independently, they must all be individually managed. In large systems, this results in a requirement for extensive network-management bandwidth for querying and otherwise managing all of the devices. With a centralized FEP approach, this would be accomplished simply by management of the FEP itself.
- The need for multiple communications connections from the device using the client-centric approach opens it up for attacks from outside. This is especially a problem because most client devices use either Microsoft Windows™ or other well-known IP stacks that have been extensively "hacked."

Given these problems, why is a client approach to providing FEP functionality considered at all? The answer is simple. It allows terminal manufacturers—be they manufacturers of client software, automatic teller machines, or terminal adapters—to quickly put a product out with minimal or no interaction with the producers of the central server application.

Server as FEP

The direct opposite to having the FEP functionality in the client is to put it all in the server. This approach is similar to that used by Tandem in the legacy world.

The primary advantage of this embedded approach is that the server application is tightly coupled to the networking functionality, simplifying management and control as well as increasing security.

The primary disadvantage of this approach is that most current applications involve multiple disciplines, each with sophisticated software applications. For example, even a simple point-of-sale terminal may use biometrics for identifying the customer as well as for doing the financial transaction. There may also be loyalty card programs, checks regarding card usage for security reasons, inventory information regarding what is purchased, management of the point-of-sale terminal itself, and, of course, management of the network links. Putting all of this in one server is, realistically, impossible. Even if achievable at some point in time, it forces enterprises to accept inferior (or no) solutions in one area in order to get superior solutions in another area. Adding new capabilities outside the server vendor's product range becomes expensive or impossible. Any attempt to move to another vendor in the future also becomes extremely expensive. Such lock-in to a single supplier is not in anyone's interests even if it is possible.

Network-Based FEP

This is the approach generally promoted by the network-equipment providers and telcos. With it, there would be a large network box at the central site that would control communications between the servers and the different network access points to which the clients would be connected.

When fully developed, this approach has the possibility of solving many of the communications problems inherent in using today's IP networks for mission-critical applications. It can provide centralized network management, a centralized audit point, separation between the network and server protocol stacks, flow control of data into the servers, and server fail-over and load balancing.

The primary disadvantages of this approach are that 1) it is not yet fully implemented in terms of the features required for large enterprises, 2) it requires all network devices (including terminal adapters) to come from a single vendor if they are to be properly managed end to end, and 3) it is purely a network device with no knowledge of either the client or server application whatsoever.

Most of the disadvantages of this approach may be resolved over time should the manufacturers see sufficient revenue from such an integrated solution to warrant the work necessary and should they also open up the management interfaces to third-party device manufacturers. However, current network boxes sold for this application are very limited and manage only the vendor's own equipment. This often leaves third-party devices operating in the "client FEP" model, the worst of all models from a mission-critical transaction application perspective.

Transaction FEP

In this approach, a box is inserted between the network and the application servers that knows of the different types of transactions that will pass through it. This box acts both as a network management and control point, and also as a message switch.

The advantages of this approach from a network perspective are the same as for the "network-based FEP." In short, it provides centralized network management, a centralized audit point, separation between the network and server protocol stacks, flow control of data into the servers, and server fail-over and load balancing. In addition, as it is transaction knowledgeable, it can also provide concentration of transactions on few communications channels to the servers, content-based message switching, and even message conversion. Finally, as it is aware of the end devices, it can provide management on a true end-to-end basis, as well as providing guarantees of message delivery not otherwise possible.

The primary disadvantage of this type of solution is that transaction FEPs are very specific to different verticals. That is, a transaction FEP for point-of-sale terminals is of little use in security, automatic teller machine, or other applications. For this reason, the spread of transaction FEPs has been slow, and most have been created by system integrators for particular customers, as opposed to being created as products for particular markets.

Summary

As can be seen, FEP functionality is vital within the IP world and will become even more important as the Internet evolves and as our commerce and personal interactions become increasingly reliant on IP technologies. Without FEPs, it will be difficult to guarantee the QoS, the scalability, or the evolution of network standards and applications that will be necessary to provide true mission-critical enterprise solutions.

The conventional approach for providing FEP functionality in the IP world is the most prone to problems, as it is based on a client-centric architecture that places the client clearly in control of all communications. Likewise, the approach of putting the functionality in the server application is not advisable, as it is based on a mainframe architecture whereby all services and applications on the mainframe were purchased from a single supplier. A network-based architecture whereby the network is managed as a single point is better but suffers from the weaknesses of being closed to third-party vendors, being incomplete, and being ignorant of the transactions that pass over it. Transaction-based FEP architectures are clearly preferable to the other approaches; however, they are not always available, as this field is just emerging.

The functionality that FEPs provide may come in a variety of guises under a variety of names (none of which are "FEP"). It is up to the purchaser to do the research and weigh the pros and cons of each approach before deciding the best approach for their mission-critical application at this point in time. This is a new era that is evolving quickly.

The Future of Communications

Brough Turner

Senior Vice President and Chief Technology Officer

NMS Communications

If the 1990s were the decade of the Web, the 2000s will be the decade of Internet-telephony communications. The industry can expect not a $10 billion change, but a $1 trillion change over the decade, with innumerable ways of marketing new voice services. But the landscape is going to be radically different from that of communications today. The key to this change is the imaginations of 10 million Web designers.

Talk about new network services dominates the industry today. Within the past 12 months, long-distance carriers have begun to admit that long-distance revenues are threatened and may no longer provide the base for company growth. The crux of the problem is that for three to four years, data bits have been overtaking traditional voice bits, but most of the revenues come from voice. The question is, what new services will generate the necessary revenues?

Recruiting Web Developers

Today, even teenagers have Web sites—a radical change from five or six years ago. And with new developments, the world of communications services is ready to open to Web developers. Imagine the possibilities if a teenager were managing communication services on a family's second phone line. There is no reason why a Web developer, even a teenager, could not come up with specialized answering treatments, download scripts, and create whatever sorts of communication services are interesting in a given week. Imagine what sounds a teenager might use for personalized alerting, with alternate sounds depending upon who is calling. For a teenager, special triggers might place a phone call to a pager or a mobile phone to announce that tickets to an event are available. Voice chat could be automatically available as part of a multiple-player game. A teen's circle of friends could have a custom find-me, follow-me intercom.

Such a scenario is not inconceivable. Remember, in five years the Web evolved from the point where most people had never heard of the Internet to where even teenagers have Web sites. With a corresponding change in the communications industry, teenagers will be able to produce their own services. No one is thinking in these terms today.

Capturing Revenues

The Internet has spawned a variety of interesting ways to market products and services. Carriers will be able to make money by capturing transactions but, in a fully evolved market, the traditional carriers may only be selling quality of service (QoS) bits. The services themselves may come from multiple diverse sources. And there will be an enormous number of new software products—some of them free.

Subscribers are still likely to contract with one company for the 7x24 execution of a session initiation protocol (SIP) script, thereby ensuring that their Internet telephone is still answered even if the power fails. That is one example of a service likely to come from traditional carriers.

There are opportunities to survive the coming commoditization of both 19th-century services—basic voice—and 20th-century custom local-area signaling services (CLASS). However, these opportunities are different from what the industry is discussing today. The voice portal business is expected to grow from literally nothing in 1999 to a more than $10 billion market in 2003. Although that is exciting growth, $10 billion is not much in terms of the telephone industry. Worldwide voice revenues are measured in the trillions of dollars.

Current Industry Focus

Today's industry focus is on softswitches and on a gradual evolution from the existing intelligent network (IN) to a "decomposed" network, which is a decomposed copy of the existing IN. In this model, there is a network operations center (NOC) or point of presence (POP) that incorporates softswitches, application servers, and media servers. There are gateways between the Internet protocol (IP) network and the public switched telephone network (PSTN), but they are all centrally located in the new equivalent of the central office (CO). Who actually develops and deploys the new applications in this environment? Traditional carriers do, exactly as in today's IN. The new application development environments are useful in that, for example, selective call forwarding can be made dramatically easier to use. However, it is still the service provider that creates the selective call-forwarding service. All the user can do is slightly modify the way the application works. The applications are still being written by the service provider or by a small group of companies that provides solutions to the service provider. The code is not available for a Web developer to play with, let alone a teenager.

Typical application programming interfaces (API), such as Java APIs for integrated networks (JAIN) and Parlay, allow third-party developers to personalize services, but most of the functionality is controlled under JAIN or Parlay. And, in this model, subscribers remain committed to a specific service provider that supplies everything. The decomposed softswitch architecture continues the model of the IN developed over the past 20 years. Although the industry is comfortable with this architecture, it is not the ultimate model.

Growth of the Internet Model

A break will occur, probably beginning within two to five years, creating the conditions for explosive growth based on an Internet model. In the Internet model, everything is wide open; nobody controls who can create a new Web site. Some of the underlying structure, such as the domain name server (DNS) and domain name ownership, is controlled, but at the application and the Web server level. No one is truly in control. That lack of central control has been very interesting in terms of the five-year transition that happened in the 1990s to the point where anyone can have a Web site.

Most of the pieces needed to build a distributed communications infrastructure that could support a transition similar to that of the Internet are in place now through the SIP family of protocols. But questions remain regarding software-development tools and adequate QoS.

The SIP family of protocols provides the foundation for distributed call control, exactly as hypertext markup language (HTML) provided everything necessary for Web sites and Web browsing. SIP does not directly cover the scripting of prompts and address other factors, but the addition of voice extensible markup language (VXML) moves the industry closer to the turning point. The SIP model supports completely distributed applications. SIP services are becoming available from major carriers around the world.

QoS Issues

Of course, the explosive growth of distributed communications services following an Internet model requires an Internet that provides "adequate" QoS for interactive voice communications. This is not available in today's public Internet. But the cost (recurring cost, not capital cost) of providing "gold" bits—i.e., two levels of service—is effectively zero when the percentage of gold bits is less than 15 percent of the total bits.

Today, data bits are growing at roughly 100 percent per year, while voice bits are growing 8 or 9 percent per year. The crossover point has already occurred, suggesting that voice will represent less than 10 percent of the bits by 2004. With zero costs and a competitive marketplace, subscribers can expect within three years to purchase Internet service for a flat monthly rate that provides today's services, plus the ability to send up to 10 percent "gold" bits for the same flat monthly fee or for an incremental flat monthly fee.

The Internet backbone already has QoS that is adequate for voice telephony. Multiple carriers offer service-level agreements (SLA) that guarantee sufficiently low latency and packet loss to support IP trunks between private branch exchanges (PBX) at different business centers, if those business centers are each connected to the same backbone provider via T1 lines. The bottleneck today is in the access network, but cable modems (CM) and digital subscriber lines (DSL) are expected to provide broadband access for 24–28 million subscribers by 2004.

A CM, today's backbone, and SIP telephones could deliver toll-quality speech service over today's Internet. QoS is an issue but not an obstacle.

Software Issues

A bigger question is that of service-creation software. For more than 20 years, service-creation environments (SCE) were dominated by high-end specialty software. The 1990s saw Microsoft's telephony application programming interface (TAPI), Java TAPI (JTAPI), S100, and a series of other APIs and environments for creating applications that had a mixture of call control and interactive voice response (IVR), with or without speech recognition. None of these really gained traction. Tens of thousands of programmers are writing applications for the communications industry, but there are three million Windows programmers and untold millions of Web programmers. The question is how to harness 10 million Web programmers and get them to write communications applications.

SIP is an early attempt to make communications call control available to Web programmers, and VXML—which has been developed in the Web portal industry and standardized by the World Wide Web Consortium—is the first step toward integrating IVR with call control. The pieces needed to make the process easy are probably less than two years away. Between SIP and various extensible markup language (XML) derivatives, over the next 12 to 24 months, everything necessary for Web developers—even teenagers—to feel comfortable creating applications will likely fall into place.

A New, Creative Approach

There is a market for services created by individuals, and when 10 million people begin thinking of services, the result will be very different from anything happening today. Such creativity will make the world of communications dramatically different.

Of course, the IN and softswitch model has years of life left because there are billions of existing phones and captive customers. The entire mobile network is built around captive services, and even third generation (3G) wireless will not change that situation in the next three to five years. On the other hand, the Internet model is set to explode, probably within the next two to five years. The transition in communications will be as dramatic as the Web transition was between 1994 and 1999.

Crossing the "Telechasm": Who Will Make It to the Other Side?

Terry Unter
Chief Operating Officer
Corvis Corporation

The telecommunications industry has reached a defining point even more profound than the digital revolution of two decades ago. The decisions carriers make today will determine whether or not they'll be in business next year.

The explosive growth of data traffic, fueled largely by the Internet, is stretching the limits of transport and processing capabilities in today's core networks. Carriers are asking how they can satisfy burgeoning and hard-to-forecast capacity demands without pricing themselves out of an intensely competitive market. If they're looking beyond mere survival, they're also wondering how to create new revenue opportunities and be first to capture them.

Data already makes up a more than half of the traffic in communications networks. Data communications continues to grow, even in a challenging business climate, doubling almost every year. These realities lead us to the conclusion that data is now the single greatest consideration in transport design. The inadvisability of applying established technologies to networks carrying traffic that is inherently changed in nature and composition seems self-evident. In the late 1990s, during the first wave of transition from voice-centric to data-centric networks, most changes in network design and technology occurred in the service-switching layer. The movement then was from circuit switching to packet switching. Now, the transition is expanding to the transport layer. The architecture and the physical building blocks for the transport layer now need to match the economics of what has become a data network. The two most significant differences between data and voice communications are the average distance of a data flow and the revenue per bit.

As carriers' traffic is increasingly dominated by data, network capacity must expand. The problem is how to expand. To continue the legacy architecture optimized for voice traffic, where the majority of calls occur within 600 to 700 kilometers, is to miss important opportunities to improve the profitability and overall performance of a carrier's business. Data communications typically span 2,500 to 3,000 kilometers. Carrier income is still primarily realized through voice services. Data services have not proved as plentiful or profitable as traditional voice. If communications growth were the result of voice traffic, the conventional design and technology of the transport layer would do nicely. But it is not, and to optimize a carrier's performance and profitability, a different kind of network is required.

So here's the turning point. Legacy wavelength division multiplexing (WDM) networks—which can carry tens or hundreds of gigabits per second per fiber—have met soaring capacity requirements, but at a price. Interim quick-fix strategies have introduced the benefits of optical-to-electrical-to-optical (O–E–O) regeneration/switching and optical add/drop multiplexers (OADMs)—technologies that enhance carrier performance up to certain capacity and service levels.

The survivors of this decade will be the ones that acknowledge that if they are successful, their networks will inevitably grow to require a new network paradigm. Ultimately, they will have to cross the "telechasm" from lost-generation to next-generation, and exploit the performance and cost advantages of optical networking.

The intelligent interim strategy, therefore, is to simultaneously position the network for today's requirements and tomorrow's realities. For any carrier that intends to prosper, the answer is to embrace *optically optimized* networks.

What Is An "Optically Optimized Network?"

The optically optimized network is one in which breakthrough optical technologies are strategically deployed where they offer the greatest return on investment (ROI) and lowest cost of ownership. The transport layer has been traditionally viewed as static capacity built to growth forecasts. As we move to a data-centric world, this transport layer becomes part of the services platform. This requires network flexibility and speed of provisioning and reconfiguration. An optically optimized network can shorten provisioning and reconfiguration intervals to days or even hours, as opposed to weeks or months. Furthermore, this flexibility allows carriers to expand their addressable markets and deliver valued services that command a price premium. Through the optically optimized network, time-of-day and day-of-week bandwidth services can be extended to customers that could not afford this capability under the old transport model.

Carriers need to be able to capitalize on the benefits of optical technologies in a way that best suits their unique network traffic, installed equipment base, service requirements, and business case. That means the specific architecture of the "optically optimized network" will be unique for each carrier:

- For many carriers, the optically optimized network is largely based on legacy equipment, complemented with hybrid electronic/optical technologies that stage for the future—with a migration path toward all-optical core networking available whenever it is required.
- As carrier traffic is increasingly dominated by data communications, the optically optimized network exploits the legacy network for local/regional grooming and transports long-distance traffic on an efficient, all-optical mesh overlay core—analogous to an express highway system

There's no question that for networks overwhelmed with data traffic, all-optical core networking is the way to go. However, judiciously applied optical innovations can optimize the performance, responsiveness, and profitability of carrier networks with lesser reach and traffic requirements—if growing their business is part of their plan.

New Network Challenges Call for New Optical Core Strategies

Internet protocol (IP)–based intranets and the Internet—and the customer demands they have spawned—have fundamentally redefined the way carrier core networks are used in two key ways.

1. Internet users demand plentiful and ever-increasing bandwidth but are reluctant to pay much for it.
2. They refuse to pay for it in any way that scales with distance.

As a result, carriers have to reconsider their strategies for handling long-haul and ultra-long-haul traffic, increasing capacity requirements and the need for flexibility in service offerings.

Unfortunately, with the type of equipment in today's networks, bandwidth is expensive and difficult to deploy, and the cost of a connection increases with transmission distance. While transport costs continue to climb, transport revenues are taking a back seat to the new revenue focus: services and content.

Some carriers have responded to this challenge with an old-world mindset: throw more legacy bandwidth and equipment at the problem. Scale instead of transform. Take what was used in the past, and install more of it. Relieve congestion in the core, and never mind that you have to manually provision multiple cross-connects and regenerators to furnish the long-haul and ultra-long-haul circuits that are increasingly in demand.

That's hardly the prescription for a networking world in which profit margins are shrinking, carriers are being called upon to provision circuits of 1,000 km or more, and profitability requires being able to offer more revenue value than simple commodity bandwidth.

"The folly of attempting to expand networks with existing SONET/SDH technologies has resulted in tremendous capital expenditures that carriers have not been able to recover with marginal revenues," states a 2001 Aberdeen Group report. "Such expenditures have led to a softening in the financial markets for carriers, and in at least one case... to outright bankruptcy."

So what *is* the prescription, if legacy optical transport isn't up to the challenge? What is the most cost-effective way to transport traffic across hundreds or thousands of kilometers? What is the most effective way to maintain signal strength over long distances and multiple nodes in the typical fiber network? How can a carrier differentiate itself in a challenging economic climate?

Available technologies offer vast differences in implementation cost, provisioning interval, performance characteristics, and total cost of ownership (TCO). Traditional network architectures can be a deck of cards poised to collapse under the weight of all their electronics. Consider a typical terabit national network, for example:

- With legacy WDM equipment, that network could occupy 700 racks of equipment at a capital cost of more than $400 million—about twice that in operating expense.
- An interim solution that introduces optical regeneration and add/drop capabilities (say, configured with optical amplifiers every 80 kilometers and O–E–O regeneration every 400 to 500 kilometers) could reasonably trim 200 racks of equipment and $8 million from the capital budget, compared to the legacy approach.
- An *optically optimized* architecture, which exploits the merits of all-optical networking where it proves in, redefines the economics altogether, as shown in *Figure 1*.

Judicious application of all-optical capabilities in the network core can completely redefine the business case for a typical terabit national network.

Is Your Network Positioned to Bridge the "Telechasm?"

How much of current network expenditure is actively returning revenues and how much is simply sustaining the network? How quickly can the network respond to changing customer requirements and new business opportunities? Can the network cost-effectively offer "liquid bandwidth," that is, poured into the network at will and served up on demand?

What is the true cost of the network—and the cost of lost opportunity? What is the cost-per-bit under the current architecture, compared to the cost-per-bit that can be achieved with forward-thinking deployment of optical innovations? That metric—comparing current to achievable efficiencies—reveals the degree to which the network is optically optimized.

How do you improve your network's score on that critical survival metric?

FIGURE 1

The Cost of Complexity

Typical 1Tbps Nationwide / Countrywide Network

	Legacy DWDM	OEO & OADM	Optically Optimized
Line Cards	4800	1600	600
Power	1982KW	800KW	304KW
Space/Racks	980	610	330
CapEx	$420M	$340M	$230M
Edge O-E-O switch ports	$90M	$90M	$90M
Core switch ports	$190M	$190M	$40M
Regens	$100M	$0M	$20M
OpEx	$800M	$600M	$85M

For Networks Not Yet Dominated by Data Traffic...
An optically optimized architecture could be achieved by implementing a hybrid "convergence switch" in the core. The Corvis optical convergence switch integrates both electronic and photonic switch fabrics for synchronous transport signal (STS)–1 and virtual channel (VC)–4 grooming and wavelength switching. In fact, it is the only product available today that combines electrical grooming for synchronous optical network (SONET) termination and optical switching in a single platform.

With this converged technology, carriers can integrate existing voice-based access, metro, regional, and national networks as well as all-optical express networks onto a single network element, with seamless network management across multiple network topologies, such as point-to-point, ring, and mesh architectures.

The carrier can then migrate traffic from SONET rings to a more efficient and cost-effective optical mesh overlay network as needed. There's no need to migrate the *network*—just migrate the *traffic* for which the business case says, "Go optical." When/if all traffic has migrated to an all-optical tier, this hybrid switch can be used in a grooming function.

This optically optimized architecture starts moving the networking away from the inherent constraints of SONET rings, particularly the requirement to reserve so much bandwidth for protection/restoration. Unlike rings, an optical mesh can be organized into many logical architectures, permitting flexible configurations that can support various traffic and protection requirements across large networks.

Multiple routes exist between network points, so traffic can be easily re-routed to an alternate path in the event of a fault or fiber cut. In contrast, ring networks would have to provide significantly more bandwidth (and support dual-ring interworking or matched nodes) in order to offer comparable protection.

For Networks Dominated by Data Traffic…
An optically optimized architecture provides an all-optical mesh overlay that serves as a high-speed core layer for express traffic (see *Figure 2*). Where legacy SONET rings are exhausting their capacity (or have proven inefficient), the carrier can migrate express traffic to the new optical mesh tier, and use freed legacy resources for collecting, multiplexing, and grooming voice and private-line traffic. The new express layer would carry the aggregated traffic delivered by electronics-based networks in addition to directly carrying the fast-growing IP traffic.

In making the decision to introduce an all-optical tier, the nature of the traffic and the growth of data communications is the primary consideration. A good rule of thumb to consider is when your network approaches 300-gigabit capacity, it is probably used heavily for data communications.

"Large long-haul carriers wishing to stay in business have no choice other than to create next-generation, overlay, all-optical express networks between NFL cities," states a November 2001 report from CIBC World Markets.

This all-optical mesh overlay network provides essential survival needs:

- Ultra-long-haul transport and all-optical capabilities reduce or eliminate the need for costly regeneration equipment.
- Open gateway interfaces accept high-speed traffic from any type of network device.
- Optical switching capability and selective add/drop functionality makes a highly flexible solution designed to drop any service at any point.
- Remote, integrated, network-management, and provisioning capabilities increase service velocity and provide for a high level of service differentiation.
- Hierarchical optical restoration greatly increases the reliability and availability of the combined network.

An effective way to deploy optical mesh backbone networks in conjunction with existing networks is to use a hierarchical, or two-tiered, architecture, where the all-optical overlay acts as the high-speed core layer for express traffic.

A Phased Migration Strategy

The good news is that it's not a leap of faith to cross the telechasm from legacy to next-generation, optically optimized architectures. A sensible migration path is available today. Carriers can start making their networks future-ready while minimizing capital expenditures.

Carriers can deploy next-generation optical transmission and switching incrementally in their existing networks and scale/upgrade as needed. With the Corvis optical network suite, optical network gateways and optical amplifiers can become OADMs. Add/drop multiplexers (ADMs) can become optical networking switches. Optical networking switches can support smooth migration from SONET rings to intelligent optical mesh architectures. New fiber routes can be added to optical switches at any point. All of it can be protected at the optical layer and managed from a unified network-management system.

For example, a carrier might do the following:

- Take the first step across the telechasm by deploying point-to-point links with optical network gateways—jumping-off points from the legacy network to the all-optical express highway
- Upgrade sites that have back-to-back optical network gateways into OADMs when optical bypass is required
- Upgrade optical ADMs to optical switches as needed

In a phased process, carriers can make forward-looking investments that 1) meet today's requirements with minimum capital expenditure, 2) position their networks for emerging demands and opportunities, and 3) enable them to upgrade whenever the business case dictates, without premature or stranded investment.

When "Optically Optimized" Means "All-Optical Mesh Core"

For carriers with capacities beyond 300 gigabits per second (Gbps) and circuits exceeding several hundred kilometers, the optically optimized architecture naturally exploits the advantages of an all-optical overlay mesh core. This architecture dramatically reduces costs, opens up new service opportunities that couldn't previously be supported, and enables the carrier to gain first-mover advantage to capture new business.

Slash the Cost of Supporting Network Traffic

"In an era when service providers, particularly competitive local exchange carriers, CLECs... are collapsing under the burden of network costs, the importance of cost containment cannot be overemphasized," states a 2001 Aberdeen Group report on the development of optical mesh architectures in the network backbone.

The capital and operating expenses of conventional point-to-point WDM systems are quite high, mainly due to the large number of regenerators required in a national-scale network. Yet in a typical backbone network, 70 percent of

FIGURE 2

A Different Approach

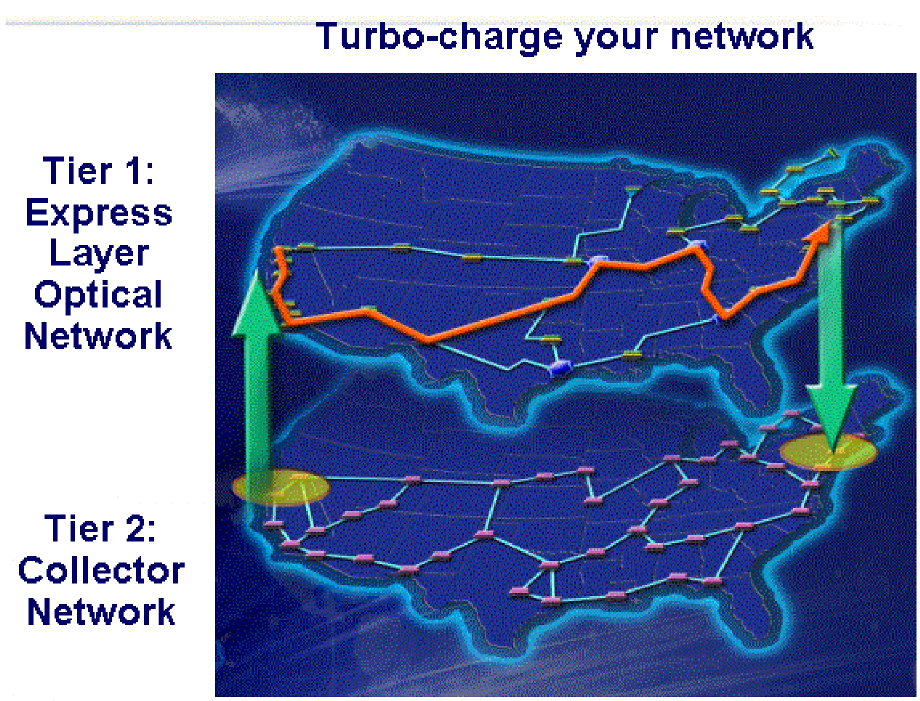

the traffic at any node is just passing through. Subjecting this traffic to unnecessary opto-electronic conversion and regeneration is a costly proposition.

"Once a carrier surpasses 500 Gbps long-haul capacity, the math proves it's time for a new long-haul architecture where the scalability of optics wins," says a November 2001 CIBC research report. If optical links of 3000 kilometers or more can be achieved without opto-electronic regeneration, some 80 percent of the circuits in typical U.S. backbone networks could do without electronic regenerators. Imagine the savings in capital expense, power, provisioning interval, operating expense, line-card inventory, and more.

Corvis systems reach more than double that distance. Its optical networking products have transmitted data more than 6,400 kilometers without electronic regeneration on a live carrier network, proving the viability of coast-to-coast express networks. A customer's "optically optimized" network may use Corvis optical products to create an overlay optical mesh core—analogous to an express highway—that minimizes or eliminates these significant cost factors:

- Number of switch ports and line cards required
- O–E–O conversions at switching points and between endpoints
- Number of discrete network-element management systems

Optical switching and transport rests on a self-evident business case: If you reduce the number of O–E–O conversion points in the network, you reduce implementation cost, network latency and complexity, provisioning interval, space and power requirements, and upgrade cost.

In fact, reducing requirements for regenerators, amplifiers, and their associated cost elements can potentially save carriers 20 to 40 percent on capital costs and 50 to 60 percent on operational cost of long-haul platforms. An all-optical overlay network takes this self-evident principle to its ultimate limit—eliminating O–E–O conversion for express traffic except at source and destination points.

Open Up New Revenue Opportunities Never before Possible

Traditionally, carriers have assessed equipment purchases from the perspective of network segments: switching, transport, budget categories, line items, engineering requirements to fulfill, etc. Viewed in that perspective, the optical express architecture is an impressive cost-saving framework, but that's only half of the picture.

Carriers need to factor in the gains that come with new service opportunities, or, conversely, the opportunities they would forfeit to competitors whose architectures enabled them to respond faster and win the business. Viewed in this perspective, the all-optical network (AON) rewrites the rulebook with instantaneous provisioning.

When you don't have to manually provision numerous regenerator sites en route, suddenly you change the economics of setting up "express" circuits (those that traverse a node without drop or regeneration), wholesaling bandwidth, and providing bandwidth on demand with point-and-click availability.

Ease of provisioning opens up the network to a whole new market segment for short-term, high-bandwidth users who cannot economically use the backbone networks as they are configured today. An all-optical mesh core creates opportunities to offer dynamic wavelength services that more closely mirror the dynamic requirements of carriers and their customers.

"This first-of-its-kind network gives us the ability to offer bandwidth that is so flexible that it can best be described as 'liquid;' which for customers means bandwidth that can be poured anywhere, anytime, in any amount," said Rick Ellenberger, Chairman of Broadwing Inc.

Rising to Opportunity Faster Than Anyone

In an intensely competitive environment with so many players, a carrier's ability to respond quickly to opportunity is a huge differentiator. It may well define which carriers survive current economic conditions in the industry.

"First mover" carriers—the ones that can turn up service in hours rather than weeks or months—will gain strategic competitive advantages in commodity bandwidth markets. Carriers that can respond this fast will do more than steal market-share; they will also uncover new customers and markets by addressing dynamic opportunities.

"The flexibility of Corvis next-generation solutions enables us to offer and deliver optical innovation to our customers in record speed."

—Matt Bross
Chief Technology Officer
Williams Communications

Optically optimized networks enable carriers to deliver revenue-generating wavelength services essentially on demand. For example, Corvis optically optimized networks support on-demand provisioning and reconfiguration. Just turn up TxRx cards at circuit endpoints; the network-management system auto-configures the rest.

As a result, carriers with optically optimized networks are setting new speed records. For example, Williams Communications provisioned an optical carrier (OC)–192 circuit within 48 hours of receiving the equipment. Broadwing Communications installed and turned up 38 OC–192 (10 Gbps) circuits in less than 60 days.

A key to this rapid delivery interval is remote point-and-click provisioning of wave services. The optically optimized network is inherently bandwidth-ready, with no requirement to preprovision expensive 10 Gbps electronics and optics at regeneration and switching sites. That means no stranded capital and no delays turning up bandwidth when and where it is needed.

"Provisioning a circuit is a keyboard exercise rather than a labor intensive, hardware driven event," said Chairman Ellenberger of Broadwing Inc. "For example, visualize having the bandwidth you need in minutes rather than months, by clicking a mouse instead of waiting for a technician."

Speed and flexibility open the door to a new world of capabilities that expand a carrier's addressable market. Carriers can now cost-effectively offer unique, short-term revenue services, such wavelength-based services by time-of-day, day-of-week, and, eventually, bandwidth on demand with tailored service-level agreements (SLAs).

Summary

As telecommunication networks are pressured to meet changing and unfavorable economic conditions, carriers are being forced to reconsider their strategies for handling traffic and services. With the type of equipment in today's networks, bandwidth is expensive and difficult to deploy, and the cost of a connection increases with transmission distance, whereas revenues do not. Legacy SONET/synchronous digital hierarchy (SDH)–based WDM networks have met soaring capacity requirements, but at high cost and with inherent constraints on scalability and efficiency.

The survivors of the next decade will be the ones that cross the "telechasm" from legacy networks to *optically optimized* networks—in which optical technologies are strategically deployed where they offer the greatest ROI and lowest cost of ownership.

For many carriers, the optically optimized network is largely based on legacy equipment for now, complemented with hybrid electronic/optical technologies that stage for the future with a clear migration path toward all-optical core networking whenever it is required.

As carriers' traffic is increasingly dominated by data, network capacity must expand. The problem is how to expand. To continue the legacy architecture optimized for voice traffic, where the majority of calls occur within 600 to 700 kilometers, is to miss important opportunities to improve the profitability and overall performance of your business. Data communications typically span 2,500 to 3,000 kilometers. To optimize performance and profitability, a different kind of network is required: an optically optimized network. The optically optimized architecture exploits the legacy network for local/regional grooming and transports long-distance traffic on an efficient, all-optical mesh overlay core—analogous to an express highway system.

With judicious implementation of optical innovations, carriers can start migrating traffic from legacy SONET/SDH networks to optically optimized core overlays and reap the benefits of optical in the time frame that suits their requirements and business case.

Carriers that adopt optically optimized networks will be able to do the following:

- Grow their businesses while controlling capital and operating expenses
- Differentiate themselves with record speed and flexibility—primarily the ability to provision new services in minutes or hours instead of months
- Expand their addressable market by offering unique, short-term revenue services and "liquid bandwidth"
- Operate with efficiencies that allow for profits despite intense competition

Enabling Multiservice Carrier Networks Using MPLS

Steve Vogelsang
Vice President and Co-Founder
Laurel Networks, Inc.

Introduction

The past several years have witnessed dramatic growth in demand for bandwidth and Internet protocol (IP) services. Service providers responded with equally dramatic investments in their core IP network infrastructure to add capacity and support IP–based services such as Internet access, transport, and Web hosting. As a result, the following has occurred:

- Service providers' IP backbone capacity exceeds that of their frame-relay and asynchronous transfer mode (ATM) service backbones, and in many cases contains unused excess capacity.
- Equipment vendors have built terabit routers that dwarf the capacity of ATM switches.

Today's economic climate is forcing service providers to slow their investment in core network capacity and instead focus on profitability. The unused excess capacity on IP backbones, along with the need to achieve profitability from their IP network investment, are causing many service providers to seek ways to derive additional service revenues from their IP network.

Service providers are now at a critical crossroad in the development of their data networks. While ATM and frame-relay networks are a key source of current revenue, their IP infrastructure represents future revenue streams and new service opportunities.

What if the excess capacity in service-provider IP networks could be used to deliver existing Layer-2 services such as frame relay and ATM? Service providers could achieve increased profitability through the following:

- Reduced infrastructure cost by deploying frame relay and ATM service over a single IP backbone.
- Increased revenue by using their IP network to support new, revenue-generating IP services.

Enter multiprotocol label switching (MPLS). Once viewed only as a traffic-engineering technology, service providers have begun to see MPLS as the missing link that can turn their IP backbone into a true multiservice infrastructure. And for good reason: Today's IP–based core routers have backplane capacity and interface speeds that surpass their ATM and frame-relay switch counterparts. MPLS brings the connection-oriented forwarding necessary for private data transport services to the IP network, while also offering the quality of service (QoS) and bandwidth-management functionality of traditional frame-relay and ATM networks.

MPLS technology combined with service edge routing enables service providers to offer existing, profitable Layer-2 private data services such as frame relay and ATM, as well as new services, such as Ethernet private line, over their IP backbone.

Multiservice over MPLS enables service providers to offer switched services as well as IP services on a common edge routing platform. Services now possible over a single IP/MPLS backbone include Ethernet private line, frame relay, and ATM services, as well as Internet access, transit, peering, and virtual private networks (VPNs).

MPLS: A Common Forwarding Plane for All Services

MPLS is an efficient tunneling mechanism with an IP control plane. It allows network engineers to create tunnels, called label-switched paths (LSPs), across their IP backbone. An MPLS–capable service edge router inserts packets into these tunnels by pre-pending one or more 4-octet labels to the packet. Interior MPLS capable devices, called label-switched routers, forward these packets based solely on these labels.

MPLS tunnels require less overhead than IP–based tunnels (such as L2TP, GRE, and IP in IP). In addition, MPLS tunnels offer the same security as frame relay or ATM permanent virtual circuits (PVCs) and are not as susceptible to denial-of-service attacks as IP–based tunnels. For example, an attacker may inject packets into an IP tunnel simply by sending traffic to the tunnel endpoint address. However, it is as difficult to inject packets into an MPLS tunnel as it is to inject them into an ATM or frame-relay PVC.

The first application of MPLS tunnels was traffic engineering, applied to the core of service-provider networks. The normal hop-by-hop, destination-based forwarding mechanism employed by IP routers tends to cause high utilization on certain network links, while other links remain relatively idle. MPLS gives service providers more control over the path IP traffic takes across the network. IP packets enter an MPLS tunnel near the edge of the network, and the service provider manipulates the path of each tunnel, enabling the provider to balance utilization on network links and increase overall efficiency.

Although traffic engineering is an important function in the network core, it is at the edge of the network that service providers can reap the greatest benefits of MPLS technology. Laurel Networks' MPLS implementation was designed specifically to support existing applications such as IP traffic engineering and Layer-3 VPNs, as well as Layer-2 transport—a promising emerging application of MPLS.

MPLS has two basic features that enable the transport of a variety of data services:

1. *Protocol Transparency:* Once an edge router forwards a set of MPLS labels, core devices forward packets based only upon these labels, not based upon any field inside the original packet. As a result, MPLS is not limited to IP transport. It can carry any type of traffic including frame relay, ATM, Ethernet, point-to-point protocol (PPP), and even time division multiplexing (TDM).
2. *Hierarchical Tunnels:* MPLS tunnels are hierarchical, meaning that an MPLS tunnel can be inserted inside another tunnel by using the MPLS label-stacking mechanism. This gives service providers tremendous flexibility and scalability, as many tunnels can be aggregated into a relatively small set of tunnels in the network core. This also has service restoration benefits, since a core network failure results in fewer connections to reroute.

MPLS: A Natural Evolution from ATM Networks

Until recently, ATM switches enjoyed a capacity and interface throughput advantage over routers, leading to widespread adoption of ATM switches at the core of service-provider networks. Some of the largest IP networks are based on a Layer-2 switched core. This design—called the overlay model—allows the provider to achieve virtual connectivity across physical backbone links, which is beneficial in terms of flexibility and traffic engineering. However, like the frame-relay switches before them, capacity-constrained ATM switches are being steadily removed from the core of most IP networks and replaced with core IP routers. The reason is speed.

The development of MPLS technology and new router interfaces based on high-speed Ethernet and packet over SONET (POS) have increased the speed and capacity of routers beyond that of Layer-2 switches, while lowering the cost per bit forwarded. Router and MPLS switch interfaces now enjoy a significant speed advantage over ATM switches. While current router ATM interfaces are limited to optical carrier (OC)–12c/synchronous transfer mode (STM)–4 (622 Mbps), POS interfaces scale to OC–192c/STM–64 (10 Gbps). ATM's cell tax—unusable bandwidth due to cell overhead—has made it an easy target for criticism. Fixed size cells allow ATM to offer deterministic delay characteristics. This functionality is most applicable to controlling latency of real-time service delivery such as voice on lower-speed access links. Yet, its fixed cell size makes ATM inefficient as a high-speed backbone.

ATM was adopted due to another advantage it had over routers: traffic engineering. Since Layer-2 switches are connection-oriented (unlike connectionless IP routers), service providers can better control the path that traffic traverses through the network to optimize network utilization and control the use of expensive long-haul links.

The need for traffic engineering in IP networks like that found in ATM networks led to the development of MPLS. MPLS provides an IP network with the same traffic-engineering advantages of the Layer-2 overlay model, while collapsing the function of routing and traffic engineering into a single device.

Benefits of an IP/MPLS Service Infrastructure

Service providers derive a number of benefits by adding Layer-2 services to their IP/MPLS networks:

- *Higher Capacity:* Demand for IP–based services continues to grow, and analysts expect revenues to surpass those from traditional frame-based services. Equipment vendors have responded by producing MPLS–capable routers that exceed the capacity of today's ATM switches.
- *Edge Service Creation:* Traditional frame-relay and ATM provisioning models require the service provider to maintain the connection state throughout the core network. In contrast, Layer-2 services over MPLS require the service provider to maintain connection state only at the edge of the network—service creation at the edge and service transparency at the core. Once inside the tunnel, all packet types are forwarded via label swapping in the core, allowing for a more scalable service backbone.
- *Common Control Plane:* Because MPLS switching is compatible with the IP routing control plane, a common control plane can be used for the converged network, improving overall network and service ease of management.
- *Fewer Connections to Manage:* In an ATM network, the call signaling rate is a critical measurement of performance, as a core network failure can cause thousands of connections to be cleared and rerouted. MPLS hierarchical tunneling requires far fewer core connections to reroute, eliminating most of the signaling burden when compared to an ATM network.
- *Network Consolidation:* A common MPLS core becomes the backbone for IP, ATM, frame-relay, and Ethernet service networks, reducing the investment required to scale existing services while enabling new services and service bundles.
- *Fast Restoration:* Advanced MPLS features such as primary/secondary explicitly routed paths, standby paths, and pre-emption allow service providers to define varying degrees of service restoration.

Combining Edge Routing with Layer-2 Services

Service edge routers combine the full IP functionality of a high-capacity edge router with the ability to create native Layer-2 services across an IP/MPLS backbone. The service edge router is an integral part of a service provider's IP network because it enables the sophisticated policy driven accounting and packet classification models needed for a provider to deliver and bill for new, premium services. The service edge router participates in Internet routing and therefore knows the ultimate destination of traffic (BGP next-hop with community), the number and capacity of the links (IGP LSDB) that the packet will traverse, and the set of ISPs (AS Path) that the packet must traverse to reach its final destination. The service edge router is capable of maintaining separate virtual routing tables for each customer site and of advertising the entire Internet routing table to a number of customers. It is this routing intelligence, combined with high-speed packet classification and accounting, that enables the use of the same IP/MPLS platform for all services.

Service edge routers are now being offered that include Layer-2 transport technology based on the following Internet drafts:

- *Transport of Layer-2 Frames over MPLS* (*draft-martini-l2circuit-trans-mpls-06.txt*) defines a signaling mechanism, based on the MPLS label-distribution protocol (LDP), which allows for the creation and status notification of Layer-2 transports.
- *Encapsulation Methods for Transport of Layer-2 Frames over MPLS* (*draft-martini-l2circuit-encap-mpls-02.txt*) defines the encapsulation for various Layer-2 frames, such as frame relay, Ethernet, ATM AAL–5, and ATM cell mode.

New Service Opportunity: Ethernet Private Line

Layer-2 transport technology (also called multiservice over MPLS) allows service providers to sustain revenue from existing ATM and frame-relay services. In addition, Layer-2 transport presents a new opportunity to extend Ethernet, a widely used and understood technology, beyond the reach of local and metropolitan-area networks (MANs). The ability to transport Layer-2 Ethernet traffic over the IP/MPLS backbone opens up the possibility of long-haul Ethernet private line (EPL) services between remote sites. This leverages the equipment already in place in the enterprise and the IP/MPLS infrastructure deployed by service providers.

The network architecture to implement EPL across an IP/MPLS backbone is illustrated in *Figure 1*. A service edge router is deployed in the service-provider point of presence (POP), delivering a Gigabit Ethernet port to the customer. The service provider may offer Internet access on one virtual local-area network (VLAN), while using other VLANs for EPL service. To deliver the EPL service, the service edge router adds two MPLS labels to each Ethernet frame received from the customer.

- The outer label or *tunnel label* defines the MPLS path between the two edge routers. All intermediate devices along the path forward packets based on the tunnel label.
- The inner label or *virtual connection (VC) label* associates the Ethernet frame with an outgoing port, enabling the service provider to support multiple EPL connections on the same tunnel.

The service provider provisions the Ethernet service by choosing two Ethernet ports to connect on the edge routers

FIGURE 1

Ethernet Private Line Service Using MPLS

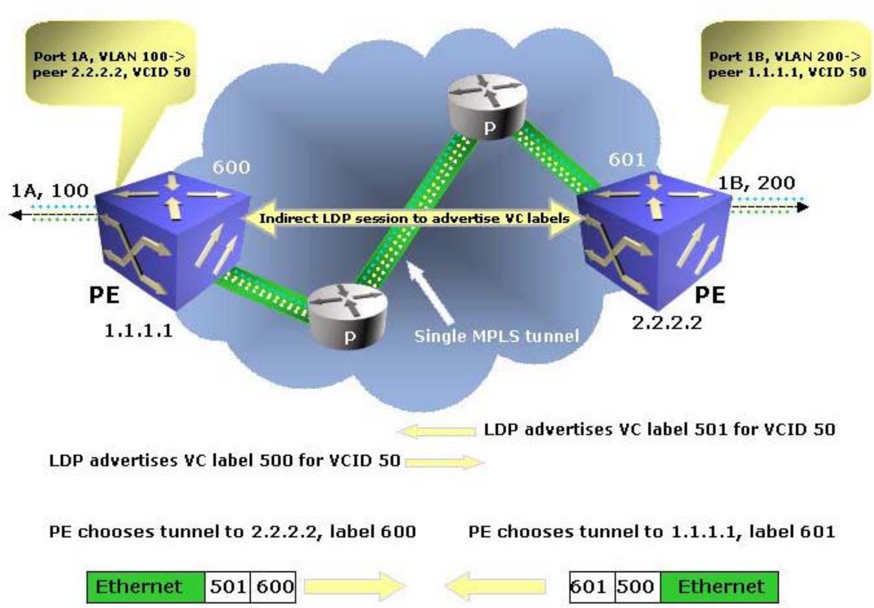

and by defining the MPLS tunnel to transport the Ethernet frames. Upon configuration, each edge router binds a locally significant VC label to its Ethernet port, and then advertises this label to its peer edge router. These labels are exchanged using LDP running in extended discovery mode, as defined in RFC 3036.

MPLS Tunnel Determines Service Characteristics

There are two choices when constructing an MPLS tunnels to carry Ethernet traffic: LDP and resource reservation protocol–traffic engineering (RSVP–TE). Tunnel creation and selection largely determine the attributes of the Layer-2 transport service offered.

Advantages of Using LDP for Tunnel Creation:

- *Simple Provisioning:* LDP can be configured to advertise only tunnels to the IP loopback addresses used to create Layer-2 service bindings.
- *Automatically Routed:* Tunnels follow the IP routed path.
- *Simple Class of Service:* Differentiated service levels are provided using simple DiffServ markings to indicate for the service class for Layer-2 transport frames on the MPLS backbone.

Advantages of Using RSVP–TE for Tunnel Creation:

- *Provides Bandwidth Reservation across the Core:* A service provider may request dedicated bandwidth for a particular tunnel, guaranteeing service levels for a group of Layer-2 transport services.
- *Provider-Controlled Path Selection:* RSVP–TE's source routing mechanism allows a provider to control the path that its Layer-2 services takes through its network (i.e., traffic engineering).
- *Optional Fast Recovery Schemes:* RSVP–TE signaling allows a service provider to control path re-optimization after a failure. Once a failure occurs and a tunnel is rerouted, the service provider may not want to automatically switch back to the previous path when it becomes available.

Conclusion

Layer-2 private data services such as frame relay and ATM are an important source of service-provider revenue today, while IP-–based services represent future revenue growth. New products combine the traffic management of today's frame-relay and ATM switches with the IP forwarding capabilities of an Internet scale router, enabling service providers to offer traditional private data services, as well new services, such as Ethernet private line, on a combined IP/MPLS network.

Driving Managed IP Service Profitability with Dense Virtual Routing

Gerald Wesel
President and Chief Executive Officer
Crescent Networks, Inc.

To build a sustainable and competitive portfolio of enterprise data services, network service providers (NSPs) need to find ways to increase the profitability of their Internet protocol (IP) services. This pursuit has led many service providers to extend beyond conventional dedicated Internet access services and begin offering a variety of managed IP services. Unlike traditional Internet access service, where a NSP simply delivers a bandwidth pipe between an enterprise user and an Internet service provider (ISP), in a managed service, a network provider takes on additional responsibilities for deploying, configuring, and managing wide-area network (WAN) switching and routing equipment for enterprise customers. Managed services enable enterprise users to reduce their operating expenses and focus on their core business by outsourcing their WAN backbone to an NSP partner. This also allow NSPs to increase the profit margins of their IP service portfolio and increase customer retention by taking over management of their enterprise customer's strategic WAN. As a result, the market research firm Gartner Group projects that managed-service revenue growth will outpace that of all other public data services through 2005, growing to a market of more than $9 billion in 2005.

This paper reviews a number of different architecture models that are available to NSPs for managed IP services, including a managed customer-premises equipment (CPE) router architecture, a network-based stacked edge router architecture, RFC 2547, and dense virtual routed networking (DVRN).

Enterprise-Based WAN Solutions

Traditionally, to interconnect their sites, enterprises deployed dedicated access routers at each of their locations and then interconnected these routers over a set of Layer-2 logical connections and Layer-1 physical links (e.g., frame relay, leased line, or Ethernet). The resulting network provided enterprises with a relatively secure backbone for their intranet traffic with consistent and predictable performance. Although the traffic of multiple customers was transported on the same physical facilities over the backbone, their traffic was separated into different timeslots, asynchronous transfer mode (ATM) virtual channels, or frame-relay connections.

In this model, the enterprise incurred the capital expense of purchasing the routers, and, more importantly, incurred the ongoing operating expense of configuring and maintaining each of the routers. Many enterprises have found it very difficult and expensive to attract and retain qualified support personnel with the necessary level of experience with routed networks.

A second approach that some enterprises have taken is to use tunneling technologies—such as IPsec, point-to-point tunneling protocol (PPTP), and Layer-2 tunneling protocol (L2TP)—to interconnect their sites over the public Internet. With this approach, businesses place a virtual private network (VPN) gateway network element and an access router at each site and configure a set of Layer-2 tunnels over a service provider's IP backbone network to interconnect their sites. All traffic is encrypted at the customer's premise by the VPN gateway before being routed to the backbone. It is then decrypted at the destination by another VPN gateway.

While this approach provides an additional level of security for enterprise users without the ongoing expense of leased lines, frame relay, or ATM connections, it requires higher up-front capital equipment expenditure for the VPN gateways. In addition, because of the complexity of configuring and maintaining the necessary tunnels and encryptions keys, this model usually results in higher recurring operating expenses than the first model.

The Introduction of Managed Services

NSPs initially entered the managed services market several years ago to enable enterprises to outsource the day-to-day management of their WAN access routers to the NSP. The

NSP combined management of CPE routers with frame relay or IP services into a single managed service, and took on the responsibility for managing enterprise customers' routers at each office location of the customers network. A number of NSPs have also begun leasing and reselling WAN access routers to enterprise customers as part of their managed service offering, giving enterprises a single point of contact and responsibility for all of their data WAN needs.

Enterprises continue to use routed networks for their WAN data traffic because it is the only way to achieve some essential requirements, including the following:

Media and Layer-2 Protocol Diversity
Routers enable an enterprise to communicate between sites with a mixture of frame relay, ATM, leased line, and Ethernet connections in their WAN.

Eliminates Need for Full N^2 Mesh of Layer-1 or Layer-2 Connections
The use of routers significantly reduces the number and cost of leased lines or virtual circuits that are required for a fully connected enterprise WAN.

Resiliency to Network Failures
Routers provide a fault-tolerant, self-healing mechanism for automatically rerouting traffic in the event of network failures.

This first generation of managed services has enabled enterprise users to save money by outsourcing their backbone WAN for data services to an NSP. However, NSPs have been prevented from scaling these types of managed CPE services to all of their enterprise customers because of the operational cost and complexity of configuring and maintaining many CPE routers in many different locations. The cost and management complexity has simply been transferred from the enterprise to the NSP.

Network-Based Solutions for Managed Services

Stacked Edge Routers
Because of the high operating expense of deploying and managing a large number of widely-distributed CPE access routers, NSPs modified their managed service offering and started to move the regional hub routers of enterprise customers to their own point-of-presence (POP) locations. This allowed the service provider to reduce the number of truck rolls required for router maintenance and the number of support staff required for their managed services. The NSP installed a separate edge router in their POP for each enterprise customer to ensure separation of route tables, support consistent performance, and make sure that problems in one customer's network (e.g., distributed denial-of-service [DDOS] attack) would not impact that of another customer. Without this type of physical separation, one enterprise customer could monopolize the memory, bandwidth, or processing resources of a router, adversely impacting the service level of other customers.

Although this network-based managed router service reduced the NSPs' operating expenses compared to the managed CPE router model, it required additional rack space and power in the NSPs' POP locations and still required a high initial capital equipment expenditure per customer and high ongoing operational expenses for configuration and management of multiple POP–based edge routers. As a result, the economics of this network-based model with discrete edge routers prevented NSPs from delivering managed services to all but their largest enterprise customers.

RFC 2547
The network working group of the Internet Engineering Task Force (IETF) has developed an approach to enable multiple IP VPNs to be delivered over a shared IP backbone network based on extensions to the border gateway protocol (BGP) routing protocol. This approach is described in RFC 2547bis, "BGP/MPLS VPNs." NSPs can use RFC 2547 to support several different enterprise customers on a single edge router, each with their own separate virtual route forwarding (VRF) instance. Multiprotocol label switching (MPLS) label-switched paths (LSPs) are also employed as tunnels to separate the traffic of different customers over an IP backbone network.

In smaller scale applications, RFC 2547 allows service providers to deliver managed services to multiple enterprise customers from a POP location without deploying and maintaining a separate edge router for each customer. However, a number of NSPs that have tested RFC 2547 and deployed it on a limited basis in their networks have identified several issues with the technology, including the following:

- RFC 2547 advertises VRF information on each customer's VPN into the global routing table, resulting in accelerated growth in the size of BGP route tables on every router in the NSPs BGP network.
- Provisioning of RFC 2547 VRFs requires extensive manual configuration with basic command line interface tools, which drives up operating expenses and lengthens the NSPs service delivery interval.
- Changes are often required in the enterprise customer's open shortest path first (OSPF) network design, since many features–such as preservation of link metrics, multiple areas, virtual links, and passive interfaces—are not supported by RFC 2547.
- Forwarding performance of edge routers decreases as more VRFs are provisioned, limiting the effective scalability of RFC 2547 networks.

RFC 2547 extends the value of service-provider investments in core and edge routers and can support multiple VPNs over a shared IP backbone network. However, the technology has still not enabled NSPs to extend managed services to the majority of their enterprise customers that have larger networks.

Dense Virtual Routing
Crescent Networks has introduced an alternative architecture for managed services called dense virtual routed networking (DVRN) that enables network service providers to deliver large numbers of dedicated virtual routed networks (VRNs) to individual customers.

The virtual router–based model, as realized with the DVRN architecture, provides features that extend beyond those in conventional RFC 2547. The virtual-router model permits each customer's service to be independently configured for routing protocols, policies, traffic engineering, and quality of service (QoS), with no impact on the customer's existing IP address scheme. The virtual-router model is the best fit for NSPs that are targeting medium to large enterprises that today run their business on leased-line, frame-relay, and ATM WANs and are looking for cost-effective managed service. Furthermore, the virtual-router model is ideally suited to complement emerging MPLS/optical networks, since it does not rely on any specific routing protocol in the core network. The DVRN architecture can run over a pure label-switched router (LSR) core or dedicated packet over SONET (PoS)/GigE on an optical backbone network.

A DVRN deployment consists of DVRN–capable service edge routers plus DVRN service-management software. The DVRN service edge router combines several types of virtual routing with MPLS label edge routing functions. Virtual-router types include virtual access routers (VARs) for individual subscribers and a virtual backbone router (VBR) for service edge router discovery, internetworking, and aggregate traffic engineering.

The VARs are then interconnected with a collection of tunnels, such as MPLS LSPs, ATM virtual circuits, or generic routing encapsulation (GRE) tunnels over an IP backbone. If MPLS LSPs are used, they may be traffic-engineered in conjunction with other core LSRs. The use of stacked tunnels provides tremendous granularity for individual customers. Stacked tunnels also enable core scalability by aggregating the many routed network instances into a manageable number of logical connections into the core. The VARs and interconnecting tunnels of an enterprise customer comprise a VRN for that enterprise over the service provider's backbone network.

To enable NSPs to scale the delivery of managed services to all of their enterprise customers, the operational cost of provisioning new customers must be significantly reduced. This can be accomplished by minimizing the amount of network engineering that must be done for each new customer and by delivering graphical provisioning and management tools that simplify and accelerate the service-delivery process.

DVRN service management automates much of the service design and provisioning process, requiring minimal intervention (through an application programming interface [API] or graphical user interface [GUI]) to customize the service for a customer. Once a VRN has been established, additional sites can be easily added to or deleted from the VRN, and additional service enhancements can be added beyond the basic managed service. For instance, a hosted storage service with specific QoS requirements can be provisioned through a single "drag-and-drop" action onto the VRN, rather than with hours of complex manual provisioning with command line interface tools. The service-management tools of DVRN can reduce operating expenses for managed services by as much as 70 percent.

The DVRN architecture is capable of supporting a wide range of applications, including managed services and conventional dedicated Internet access services. For instance, an NSP can initially offer a standard dedicated Internet access service, extend the service with a managed service offering for improved enterprise connectivity, and then bundle the managed service with xSP infrastructure services, such as content delivery, hosted storage, and applications hosting. Each VRN operates as a distinct customized routed network tuned to the needs of its particular customer and applications. A NSP may set up VRNs to seamlessly connect xSPs with their enterprise customers. VRNs can also be provisioned to provide xSPs a way to expand into new geographic regions.

DVRN provides a platform for a cost-effective and seamless suite of managed IP services—such as intranet and extranet—along with conventional dedicated Internet access services. VRNs can create dynamic any-to-any connectivity between a corporate headquarters over optical Ethernet or leased-line connections, branch offices over frame relay and digital subscriber line (DSL), teleworkers, and partners over remote-access VPN tunnels. If an enterprise has already deployed a routed WAN, this network can be outsourced to a provider-based managed service with no impact on the customer's local-area network (LAN) or routed backbone network architecture.

The DVRN architecture delivers several significant advantages other alternative approaches for managed services, including the following:

- Significantly lower capital expense than approaches based on stacked edge routers or RFC 2547
- Dramatically reduced operations expense based on use of traffic-engineering protocols and graphical provisioning and management tools
- Protocol flexibility—DVRN enables service providers to offer their customers their choice of interior gateway protocols (IGPs), along with supporting managed services over ATM, IP, and MPLS backbone networks
- Enhanced scalability, including the ability to support hundreds of enterprise and service-provider customers from a single platform

As a result, the use of DVRN can dramatically increase the profitability of NSP managed IP services over alternative network-based or enterprise-based solutions.

Summary

In their efforts to improve the profitability of their IP service portfolio, NSPs have begun to offer a variety of managed IP services to their enterprise customers, allowing these customers to outsource the deployment and management of their data backbone network to the NSP. Several architectural models have developed for the deployment of managed services, including a managed CPE router, stacked edge router, RFC 2547, and DVRN. The DVRN architecture can dramatically increase the profitability of NSP managed IP services by lowering the initial capital equipment expense, and ongoing operating expense, of managed services.

Broadband and the Opportunity Before Us

Edward E. Whitacre, Jr.
Chairman and Chief Executive Officer
SBC Communications Inc.

My discussion here focuses on the need for balanced broadband regulation. It reminds me of the time a young teacher was up for a job in a small, unsophisticated country town. During the interview, the principal asked the young man if he taught that the Earth was round or flat. Anxious to get the job, he said, "Well, I can teach it either way."

When it comes to talking about broadband, I'm sure many of us feel the same way. I know where our competitors stand on the issues. I understand most regulators concerns. And I'm sure you have a pretty good sense of my views. We are familiar with each other's positions, and we know what familiarity breeds, right? However, I want to say at the outset that there is no room for contempt in this issue. I appreciate the stake that everyone has taken in this important debate. I have done my level best to understand them. Of course, there are some legitimate disagreements.

But as I look over the telecom landscape, especially broadband, I see that we share many common goals. And I see a lot of room for agreement. I think all of us are optimistic over the prospect of a fresh start in 2002, particularly after the year we went through in 2001.

But we have more to be hopeful for than just a fresh start. I believe that there are significant opportunities just within our grasp. Opportunities to provide Americans with an amazing array of new communications products and services. Opportunities to create jobs in an economy that lost 2 million last year—more than 300,000 in the telecom sector alone. Opportunities to pull the nation's contracting economy out of the doldrums by generating billions of dollars of growth annually. I believe that we have before us the opportunity to return the information technology (IT) sector to its role as one of the leading elements of economic expansion in America—and to ensure that Americans remain at the forefront of the technology curve.

We can do all of that if we, as industry leaders and as policymakers, accept our responsibility to get the policy regarding broadband right—and if we act quickly.

I mentioned that I see many of areas of agreement on this, and I would like to spend a little time now reviewing those areas. If we can accept this set of facts, I believe it will lead us all to a similar conclusion about the need for a consistent national broadband policy.

First, we share a common goal of deploying broadband to consumers and the businesses. There is remarkable agreement on this point—from Silicon Valley to the marble hallways of Washington. I do not think that anyone seriously disputes the importance of broadband. Certainly, everyone understands the value of bigger, faster connections to homes and businesses. We recognize the incredible transformational potential of broadband in communications, in entertainment, and in how we do business.

Brookings Institution economist Robert Crandall estimates that accelerated broadband deployment could generate half a trillion dollars annually in economic growth. Now *that* is an economic stimulus package worthy of the name. And the even better news is that these benefits would come with no strings attached. No increased government spending, no taxes—nothing—so long as we allow the market to work fairly.

The second issue that we agree on is that consumers will certainly benefit from a wide array of choices. Consumers are best served when many competitors offer services over their own facilities using their own technologies. The good news is that situation exists today—for now. Today, there are at least four technologies competing to offer essentially similar broadband services: cable modem, satellite, fixed wireless, and digital subscriber line (DSL). None controls the only way into a consumer's home or place of business. Each serves its customers over its own facilities.

However, regulators have so far focused only on mandating additional competition for DSL service, but consumers do not distinguish among technology choices. Consumers only see a single broadband market, which is how it ought to be regulated.

How attractive is that market? The evidence is everywhere. It is in the $72 billion that Comcast is paying to acquire AT&T Broadband. Once that sale is complete, Comcast will become the nation's biggest cable company with 22 million subscribers. Not surprisingly, it will be the biggest broadband provider, too. The evidence is also in the $26 billion that EchoStar is paying to acquire DirecTV. Once that deal closes, EchosStar will command the satellite market, with a customer base of more than 16 million subscribers. And the evidence is in the billions that SBC is investing in new networks to bring DSL to our customers, and the billions more that other telephone companies are spending as well.

The third issue we can agree on is that all of this potential is tempered by reality—and the harsh reality is that for a variety of reasons, including uncertainty in the regulatory environment, America's broadband leadership is in question and in danger of slipping. Today, less than 10 percent of Americans have broadband service. That percentage is simply too low to attract the investment needed from software developers and equipment manufacturers to make the promise of broadband a reality. We have not yet reached the critical mass needed to produce the kinds of services that broadband can and will one day provide.

Worse still, we are falling behind other nations in the race to bring broadband to the mass market. Growth of broadband in the United States is significantly behind the growth rate in other regions. Take DSL, for example. During the second quarter of 2001, DSL grew at a rate of 35 percent in Europe and 37 percent in the Asian-Pacific. Here in America, DSL grew about half as fast in the same quarter at about 15 percent.

No one wants to see this nation lag the world in broadband. America cannot lead the world in growth if we do not act soon to remove the narrowband governor on our economy. In my view, these are the broad issues on which we can all agree: The telecom and broadband industry is hurting right now, which stymies the delivery of advanced services to consumers.

The potential of the broadband marketplace remains huge, as well as attractive enough to competitors that we are making equally huge bets on it. And the market is open to a range of companies, competing via different technologies. So, the logical question to ask is "What's the hold up?" We have the ingredients for a perfect storm of broadband competition: sizeable customer demand, numerous companies offering competing technologies and willing to invest big, and no barriers to customers. So why aren't we seeing more rapid deployment of broadband?

For us, the answer lies in one last indisputable issue: Of the competing broadband technologies, only one is regulated, and pervasively so. And in a strange twist, the regulated provider is not the one that controls 70 percent of the market. It is not the one that is closed to competitors that want access to their networks. That, of course, is cable.

Instead, it is telephone-company–provided DSL service that is regulated. It is DSL, which serves about 28 percent of the broadband market, that has been subject to rules and regulations intended for a bygone era of communications—while cable threatens to run away with the market, to the detriment of consumers.

The Telecommunications Act of 1996 set up rules under which our voice networks were opened to competitors: co-location, line sharing, deeply discounted pricing. We accept that for our legacy voice networks, which were built when we had an exclusive franchise, at prices set to recover costs. Yet this strategy does not work for new investments in new networks, where no cost recovery or customer is promised.

By applying yesterday's rules to today's networks, we have, in George Gilder's memorable phrase, "privatized the risk and socialized the reward." For incumbent local-exchange carriers (ILECs), such as SBC, the message is clear: Make the investment in new networks. Take the financial risk. Carry the excess capacity. If the bet pays off, competitors can piggyback on that investment at rent-controlled rates without investing a dime and without deploying any new technology of their own. If the bet fails, our shareholders are left holding the bag.

This is not an academic argument. The environment that we are in has real-world affects on our ability to deploy new networks efficiently and on our ability to compete effectively. Here's what I mean: In 1999, we unveiled a $6 billion project to bring broadband to many more customers across our region. But first, we were required to create a brand-new, standalone data subsidiary to provide DSL service. We had to hire, train, and transfer new employees; create new databases; transfer the network; and transfer customers to the new company. We created complex, duplicative systems that by themselves added $250 million in costs to us but zero to the value of the service being provided. You can imagine the millions it cost to create a new subsidiary, but you cannot imagine the damage it cost in terms of our ability to efficiently meet customer expectations. The cable companies cannot imagine it either, since they face no similar requirement.

Still, we plowed ahead with our new network—key elements of which are the thousands of neighborhood broadband gateways that we will install. These units house the electronics that make DSL possible. We are putting them in neighborhoods across our territory—rich and poor, urban and suburban—without cherry-picking communities. However, before we could turn these broadband gateways on, we asked for a ruling on a simple question: Who should own the equipment in the gateways—our telephone company or our advanced services subsidiary? The answer would determine how competitors would use our equipment.

This took almost nine months to resolve, thanks mostly to our competitors, which turned it into a contested regulatory proceeding. In that time, unregulated cable companies raced ahead with their deployment. We lost ground to them. And millions of consumers were forced to continue waiting for DSL.

Finally, the issues were settled with a great, big "yes…but." *Yes*, we could move forward, *but* SBC—and only SBC—was required to make our broadband gateways bigger to accommodate potential competitors. We were mandated to deploy optical concentration devices and to implement other requirements. All of this raised the cost of DSL deployment by more than $350 million, and here's the kicker: No competitor is using the optical concentration devices (OCDs) or the remote terminals (RTs). None.

Competitors may have gained something from the lengthy delay, but consumers and DSL technology were the losers. It is a simple proposition. If higher costs are incurred, customers must pay more for the service. Or, they will choose another technology instead. It diminishes our ability to deploy DSL, and the technology is harmed.

Broadband deployment requires massive capital investment. Responsible companies cannot make this level of investment on a whim. They do not do so if they believe the market is closed or if the opportunity to effectively compete is shut off. At SBC, we are investing in new networks because we are convinced of the viability of the broadband market, and we want to be a leader in it. But the reality is that the capital risk is squarely on our shareholders. While there is good demand for broadband, no one is promised anything. Some people do not use the Internet now and never will. Others are satisfied with dial-up. And those who want broadband have an array of choices.

However, if regulatory burdens cause DSL to be priced out of the market, we must make rational decisions and adjust our investment accordingly. This affects competing competitive local-exchange carriers (CLECs) as well as customers, but customers stand to lose the most if cable is left as the dominant, unregulated provider. That's not an outcome that anyone wants.

This asymmetrical regulation will lead us to an outcome that is bad for consumers and for our economy. It creates a disincentive for companies such as SBC to invest in and deploy new advanced networks. It forces the provider best able to compete against cable to slow down and reassess its strategy instead of pursuing it as vigorously as possible. That is certainly the case at SBC, where we have had to slow our broadband deployment.

The asymmetrical regulation also gives all competitors but one free rein to operate without concern for added cost and delay that regulation inevitably brings with it. That's bad news for my sector of the industry, but it's worse news for consumers who deserve a full range of broadband choices. On top of this, our competitors are urging state regulators to consider even more onerous and more expensive requirements on broadband, further clouding the future with uncertainty. So I hope it's clear why we so strongly support a coherent, pre-emptive federal broadband policy.

For us to seize the opportunity to keep American consumers at the vanguard of the broadband revolution, policymakers must seize this moment and let the free market work freely. We must create a consistent national policy on broadband and other advanced services. SBC and others like us want to compete under the same rules and investment incentives that apply to the dominant player in the broadband market. The real beauty is that removing the disincentives that we face would cost nothing—in terms of money or in the ability of competitors to maintain access to our networks. Competitors would keep their ability to resell our service or use our network to reach their customers. We will continue to make our networks open to Internet service providers (ISPs) without discrimination. In fact, by leveling the playing field, the only real threat of discrimination that hangs over this debate would be eliminated. Singling out one non-dominant technology for regulation, while the others are left alone, is unfair.

Common sense tells us that by allowing this asymmetrical environment to exist, policymakers have put themselves in a tough spot. They may not be perceived as having picked the winner in the broadband race, but they could be viewed as having settled on a loser. I know that's not what any of us wants. And I do not believe that is anyone's intent. But that is the potential affect.

When you get right down to it, the only threat is in not taking action. The threat lies in maintaining the status quo or allowing it to deteriorate into 50 different sets of rules and regulations. A patchwork quilt of regulations that apply to only one competitor—the one that is a distant second to the market leader—will not produce the results that this nation needs.

Yet it does not have to be this way, as history shows. Industry and policymakers have faced similar situations before, and made the right decision, to the benefit of millions of consumers. Years ago, the Federal Communications Commission (FCC) decided against applying legacy voice network rules to the emerging wireless industry. It deregulated the industry and pre-empted the states. Artificial long-distance boundaries were erased. Policymakers took a hands-off approach, choosing instead to let the emerging industry develop and mature on its own. No regulatory burdens were added. No disincentives were created. Many competitors came into the business and invested heavily in their own technologies and networks. As a result, the industry has flourished. Prices continue to fall, while innovation proceeds swiftly.

We can produce the same results in broadband. And I am optimistic that we will—not just because it's the New Year, but also because of some very encouraging words from FCC Chairman Michael Powell. Last October, Chairman Powell made it very clear that broadband infrastructure deployment is the FCC's central communications policy objective. In his announcement, he said that "substantial investment is required to build out the networks, and we should limit regulatory costs and regulatory uncertainty." And he added: "Broadband service should exist in a minimally regulated space."

When I heard those remarks, it was one of those "I wish I'd said that" moments. Still, we all know that stating policy goals and achieving them are two different things. Yet I believe that this is one goal that must be met. The promise that broadband offers to our economic well being and technological leadership is real. And the opportunity is now.

We can do exactly what Chairman Powell has described, in a way that protects competition and ensures a level playing field for all. Accomplishing that goal would be a resolution that every American would benefit from seeing us keep.

This presentation was made on January 14, 2002, at the Emerging Issues Policy Forum.

Integrating Enhanced Voice Collaboration Services in a Converged Network

Ed Yackey
Chief Technology Officer
Voyant Technologies

This paper will discuss how the network is changing and evolving, how those changes are reshaping voice collaboration, and what the future of voice collaboration will be within these networks.

The Network Is Shifting from Monolithic to Distributed

A recurring theme in today's communications market is the shift from monolithic to distributed networks. This evolution has affected the overall development and deployment of applications in the network.

The Monolithic Network

For example, the traditional circuit-switched architecture uses a monolithic network. It has a mainframe mentality, where one vendor provides all the services and equipment. It tends to be a closed and slow-moving environment in which it is difficult to drive innovation and be creative. Its advantages are that it is well understood and very stable. Additionally, from a service provider's perspective, an advantage is that a single accountable vendor provides a single point of contact when there are problems.

The Distributed Network

As the distributed network emerges, Internet protocol (IP) is becoming the foundation of the emerging packet-switched architecture. The distributed network offers integrated services and best-of-breed components. It is an open environment with fast time to market for new applications. Its main disadvantage is the uncertainty in standards that are still evolving. This has lead to integration challenges.

In the distributed network, services are "in the cloud." These services may include presence, voice conferencing, data conferencing, voice recognition, and billing services. Services may be interdependent or can be tied together to create new applications. Here, the focus is on the application, not the technology.

Some advantages of a distributed architecture include the option to select best-of-breed services, the ability to add new services without affecting existing ones, the capacity to scale each service independently, and the ability for multiple applications to use the same services. The distributed architecture significantly lowers the barriers to developing and deploying new applications, and it allows specialized vertical applications.

The Future of Collaborative Voice

The shift from a monolithic to a distributed network will change voice collaboration in several ways. It will make collaboration at a distance more convenient by making it possible to connect from any communication tool. It will make collaborative voice easy to use by allowing meetings to be controlled in a natural manner—not via touchtones. Finally, it will make voice collaboration more cost-efficient, removing price as a barrier.

These effects could result from the distributed network providing various endpoints, such as wireless, IP, personal computer (PC)–based software phones, and public switched telephone network (PSTN) phones. The network will also allow versatile controls, including controls based on voice, touchtone, PC, personal digital assistant (PDA), and wireless phone "browsers." Seamless convergence will make services natural and easy to use, hiding underlying technology. The network will address cost issues by leveraging technology to reduce operation, network, and equipment costs.

Network Innovation + Ease of Use + Application Adoption = Growth

Reservation-Based Audioconferencing

Around 20 years ago, audioconferencing was first created and deployed on reservation-based systems. In a reservation-based system, audioconference calls are considered events. The call has to be planned ahead and requires operator assistance. Back then, conferencing resources were limited and required a large operations staff. Thus, conferencing was treated as a very scarce resource.

Subscription-Based Audioconferencing

In the mid-to-late 1990s, subscription-based models for audioconferencing began to emerge. In this model, a conference user establishes a financial relationship with the service provider upon subscribing to the service. Users are given a phone number to use whenever they want to have a conference. Whenever that phone number is dialed, the conference is available and ready to use. Along with that subscription comes a profile that specifies the conference; for example, how many people might be in a conference, whether international dialing is included, whether it is possible to dial out, etc. The features associated with the profile are stored by the system.

In moving from a reservation-based model to a subscription-based model, existing customer usage typically increases 20 percent. In some cases this increase is seen almost immediately. This jump in usage exists primarily because of the system's ease of use and because an operator is not required for a conference. But subscription-based conferencing still accounts for only about 1 percent of long-distance telephony minutes. Thus, subscription-based conferencing is only a step in the overall evolution of collaborative voice.

Instant Conferencing

The next step is the instant model, in which the system itself has no forward knowledge of the event. Removing the need to provision customers in the system removes all ties from the underlying services to the application itself. This allows more applications to use the same services.

Figure 1 shows a graph that compares ease of use and volume of calls for collaborative voice services over time. After more than 15 years with no significant growth in reservation-based systems, growth in reservationless conferencing skyrocketed in only five years. Following that growth, the instant model, which uses IP, is expected to take hold as market trends continue. Ease of use, user interfaces, control, time, convenience, and transparency are all improving. There is an endpoint explosion, driven by wireless technology. Access to additional resources is also increasing, and the price per minute continues to decline.

A Transformation toward Unified Communications

Another way to look at this point is shown in *Figure 2*. Voice collaboration started as a stand-alone application with a decent-sized market, but it was not that interesting by itself.

As other applications were added along with voice conferencing, service providers started to offer voice- and data-sharing applications. Today, several service providers offer unified applications with voice conferencing and data conferencing. The applications are deployed to the end user as one service. End users are unaware that two systems underlie this application. Users receive access to a Web page that allows them to control both their data conferencing and their audioconferencing.

As the industry moves forward, audioconferencing is becoming a smaller piece of a larger application pie. It is becoming more of an enabling technology that other applications use. This is true for voice conferencing and for other services as well.

Example: IM + Presence + Conference

The following example shows how some of these services can start to be deployed in converged networks. Gary, who is at his desk, wants to talk with three coworkers. Typically,

FIGURE 1

Network Innovation + Ease of Use + Application Adoption = Growth

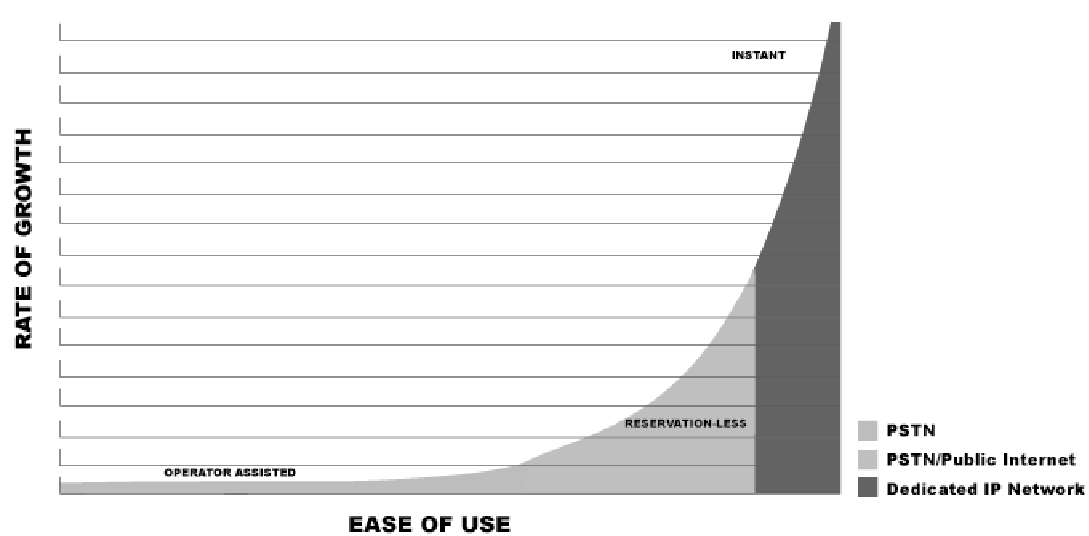

FIGURE 2

A Transformation toward Unified Communications

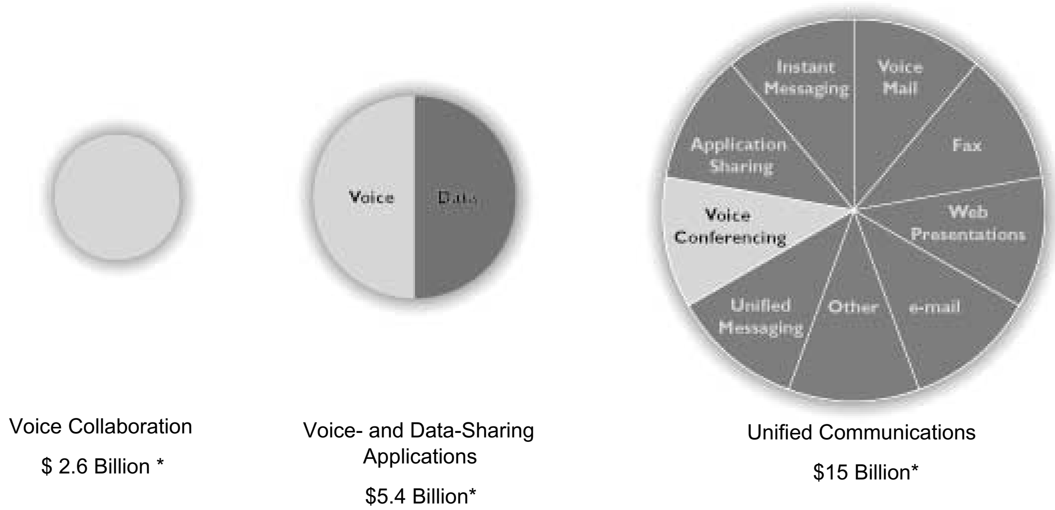

* Revenue figures are for 2003; cited by Telespan, IDC, and Piper Jaffery, respectively

he might walk outside his office to try to find the coworkers and pull them into the conference room. However, in this example he looks at his desk to check their presence and sees that none is currently available. He then logs onto his presence server on his desktop, clicks on the three people he wants to talk to on the PC instant messaging (IM) client, and then clicks on "voice" and "as soon as possible." That request registers with an application server in the network.

When all four people are available, the system initiates the voice conference by dialing out to them and placing them into a conference call. The people at their desktops communicate over their PSTN phones. Others talk over IP or wireless phones. If there is a need to share a slide or a spreadsheet during the conference, instead of stopping and restarting things, Gary can simply select "data" on his IM client to start a data conference and push up a slide show or spreadsheet. The slide show or spreadsheet will be shared over a PC at individual desktops, or over a PDA for users on the road.

Coupling the presence concept with the conferencing concept provides all the power needed to schedule the conference. Gary, in the example, does not have to know how to reach each person. All he has to do is enter a request to talk to them. Convergence of these different technologies and networks starts to provide a very rich experience that affects the end user's business day in a positive way.

Manage the Internal Changes to Fit the External Changes to Win the Competition in the New Era

Robert Zhang
Chief Executive Officer and President
ZoMatch (China), Inc.

The wireless telecommunications industry is changing rapidly. Because of the inventions and applications of new technologies, the new business model that is suitable to the new technologies and applications needs to be developed. Therefore, it is very critical how a company manages the internal changes to fit the external changes through aligning strategy and organization to position itself to win the competition in the future wireless telecommunications industry, especially the radio frequency (RF) components industry.

Background

Today' telecommunications industry is being shaped by several key factors: a rapid growth in mobile communications, explosive growth in broadband, stagnating revenues from core telephone services, falling prices and spreading competition, liberalization, and privatization.

Present and Future External Environment

According to research companies' studies, wireless mobile business is continuously growing strongly. In the developing market, the continuous expansion of the mobile network is necessary to support subscription growth. While in the developed market, the demand for mobile Internet is under the development. Third generation (3G) is close to implementation. Wireless data also looks as though it is going to be a huge business, such as IEEE802.11x standard–based wireless local-area network (WLAN) and IEEE802.16 standard–based wireless metropolitan-area network (WMAN) applications. According to J.P. Morgan's estimation, the European market alone for wireless data will be worth some $82.4 billion by 2010, with wireless data revenues in 2006 being larger than the existing voice market (by revenues) in 1999.

Furthermore, the three big wireless markets—the United States, Europe, and Asia (mainly China, Japan, Korea, and Taiwan)—are all growing, and the Chinese market is growing most rapidly. For example, looking at the Chinese market situation, the subscriber base increases at an unbelievable rate—the total for mobile subscribers there was about 140 million at the end of 2001. However, the penetration rate is only about 10 percent of the total population. There is still a lot of room to grow. More wireless infrastructure needs to be deployed to either increase network capacity or cover the uncovered areas. This is good news for equipment suppliers.

All of those factors are pushing up infrastructure demand. Similarly, mobile infrastructure will be heavily invested in during the next few years. However, on the other hand, the price erosion of telecom equipment is becoming worse. Because of this, the demand for RF components, with the lowest price for a given performance requirement product, will be huge and will continue to grow for quite a few years. The RF components business is not the same as the wireless or mobile infrastructure business. It doesn't depend on protocols. No matter what types of protocols the system houses adopt, they will need to have RF components to build their system.

RF Component Market Characteristics and Strategic Considerations

The RF component market is in high speed, characterized by rapid-fire technological change, short product cycles, and rapidly evolving customer requirements and expectations—all occurring at once. To be competitively successful in fast-changing markets, the following elements need to be considered when developing a corporate strategy:

- Invest aggressively in R&D to stay on the leading edge of technology know-how. Translate the advances into innovative new products. It is also important to focus R&D efforts in a few critical areas to avoid stretching the limited resources too thinly and to deepen expertise.
- Develop the organizational capability to respond quickly to important new events. Quick reaction times are essential because it is impossible to predict or foresee all of the changes that will occur. Resource flexibility tends to be a key successful factor in creating new

competencies and capabilities. Being a fast follower, if not the first mover, is critical.
- Rely on strategic partnerships with outside suppliers and with companies making tie-in products to perform those activities in the total industry value chain where they have specialized expertise and capabilities.
- Cutting-edge know-how and first-to-market capabilities are very valuable competitive assets.
- Strike a good balance between building a rich internal resource base that, on one hand, keeps time from being at the mercy of its suppliers and allies and, on the other hand, maintains organizational agility by relying on the resources and expertise of outsiders.

Strategy to Win the Competition

Companies need to develop suitable strategies according to external and internal environment and conditions. More importantly, the strategy that a company is going to develop or adopt needs to be able to utilize one's core competencies, create competitiveness, and build the entry barrier for other competitors. Because of the nature of the RF component industry, the strategy of being a best-cost supplier is recommended.

The reason to adopt this strategy is that it aims at giving customers more value for their money. It combines a strategic emphasis on low cost with a strategic emphasis on more than minimally acceptable quality, service, features, and performance. The idea is to create superior value by meeting or exceeding customers' expectations of key quality/service/features/performance attributes and by beating their expectations on price. The aim is to become the low-cost provider of a product or service with good-to-excellent attributes, and then use the cost advantage to under-price brands with comparable attributes.

What distinguishes a successful best-cost provider is having the resources, know-how, and capabilities to incorporate upscale product attributes at a low cost.

The advantages of this strategy are as follows:

- It has great appeal from the standpoint of competitive positioning.
- It produces superior customer value by balancing a strategic emphasis on low cost against a strategic emphasis on differentiation.
- It is a hybrid strategy that allows us to combine the competitive advantage of both low cost and differentiation to deliver superior customer value.

In markets where customer diversity makes product differentiation the norm, and where many customers are price- and value-sensitive, a best-cost producer strategy can be more advantageous than either a pure low-cost producer strategy or a pure differentiation strategy keyed to product superiority. This is because a best-cost supplier/provider can position itself near the middle of the market with either a medium-quality product at a below-average price, or a very good product at a medium price. Often the majority of customers prefer a mid-range product rather than the cheap, basic product of a low-cost producer or the expensive product of a top-of-the-line differentiator.

The next question is how a company can realize this crafted strategy—i.e., how does one implement it? The recommended strategy suggests that we need to develop an efficient organization to implement it. Because the organization consists of processes and individuals, we thus need to look at all of these levels.

Organization to Align with the Strategy

To build an efficient organization to fulfill the aforementioned requirements, one should focus on developing the organization and improving the performance from three levels: organization, process, and job/performer. With each level, three factors/variables will be adopted to evaluate performance needs, which determine effectiveness at each level and the effectiveness of any system. Combining theses three levels with the performance needs results in nine performance variables. *Table 1* illustrates the comprehensive set of improvement levers that a company's managers are suggested to use at any level.

To implement and apply this, one should do the following:

At the organization level, design the organization to implement the strategy effectively and efficiently—i.e., organization structure. Key success indicators need be developed throughout the targeted implementation areas in the company.

At the process level, key processes/procedures need to be developed, especially regarding manufacturing and quality

TABLE 1

	Goals	***Design***	***Management***
Organization Level	Organization Goals	Organization Design	Organization Management
Process Level	Process Goals	Process Design	Process Management
Job/Performer Level	Job Goals	Job Design	Job Management

control. Here is a sample list of processes/procedures needs to be developed:

- Business decision-making process
- Risk-control process
- Go/no-go business decision-making process
- Marketing process
- Customer-assessment process
- Resource -allocation process
- Project-management process
- R&D process/procedure
- Manufacturing process/procedures
- Quality-control process
- Supply-chain–control process

At the individual level, design the value and activity chains and jobs. Incorporate a system of reward and recognition that promotes implementation of the plans and stimulates employees' motivation. Design jobs, select people, and build and focus competence on corresponding key activities. Set team and individual performance objectives and allocate resources and support for achievements.

- Set progress milestones and review dates using project-management and evaluation tools to improve performance and measure progress.
- Identify major business risks and uncertainties. Develop control-driven contingency plans and protocols to deal with the unexpected.

Plan for Change

Today's telecommunication industry is an emerging industry. The speed of telecom-industry development is very fast compared with other industries. The external environment and competition landscape change quickly and frequently. The demand for organizational change has accelerated at an extraordinary pace in recent years. For example, mobile technology (from analogue to digital) has not only opened a big market for voice/data communications, but it also is replacing traditional circuit-based technology in the voice market. In the near future, the broadband wireless integrated services (voice, data, and image) will further impact the services on fixed wireline networks.

Enormous industry changes have brought us breathtaking opportunities. However, due to the needs to adopt new technologies and new business models, the old business strategy and organizational structure are under great pressure to be changed or modified. The companies in the telecommunications industry are all facing questions such as the following: What does it take for a company to be a part of this dynamic industry? How can we develop a strategy and build up an organizational structure that can adopt the external changes? This is indeed a challenge.

Modified Integrated Strategic Change Model

Considering the continuous changes in the outside environment, a Modified Integrated Strategic Change Model (MISCM) is suggested for adoption in order to fully accommodate external environment changes (see *Figure 1*).

While implementing strategic change, a company has to examine the alignment of strategic change. For the strategic analysis period, one needs to search for misfits, to pay attention to the process of aligning, and to define the specification of future alignment.

FIGURE 1

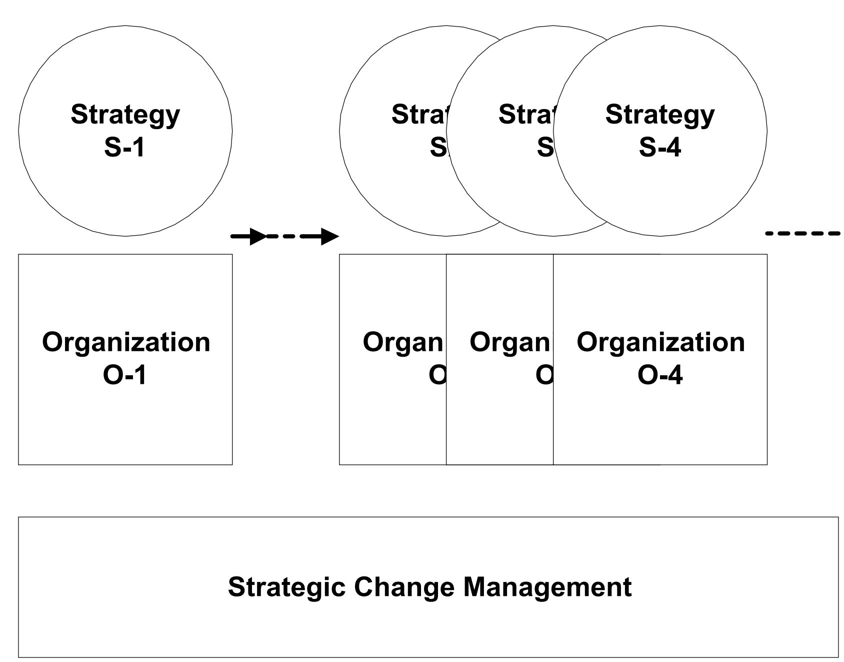

The MISCM suggests building up a timely strategic change-management system in the organization to manage the internal changes (strategy and organization) to fit the external changes. The MISCM provides a long-term source of competitive advantage—the capability to create and implement strategic change over and over again. Understanding this objective requires an understanding the sources, nature, and limits of competitive advantage. In the long run, however, any source of competitive advantage, while salient and important for a period of time, eventually shifts to other sources. The resources that once conferred an advantage to a company become irrelevant or depreciate in value. Thus, for any company, even sustainable advantages may provide performance enhancements only for a given time period.

Summary

Although the telecom industry is changing rapidly and many telecom companies failed quickly, I believe that a company is able to sustain its success as long as they successfully manage the internal changes (strategy and organization) to fit the external changes. It is a key management factor to being successful in the future, even faster-changing, telecom industry.

Section II:

Business, Marketing, and Regulatory Issues

Telecommunications Industry Overview

Market Dynamics Are Driving the Transition to a Converged Voice and Data Access World

Rob Avery

Vice President, Marketing

Gluon Networks, Inc.

Overview

We are living on the edge of the digital divide. Virtually every telephone worldwide is connected through an obsolete access network to a local switch that is not capable of delivering the services that users demand. Internet content and corporate virtual private networks (VPNs) have become richer, while large investments have been made in the core network adding up to terabits of capacity.

The access network and central office (CO) infrastructure are collectively viewed as the last remaining bottleneck to universal enhanced services. Two major forces are driving these changes in the telecommunications industry today: deregulation and technology.

Deregulation is driving the exponential growth of competitive services. The break-up of the Bell System in 1984 and the Telecommunications Act of 1996, which resulted in the creation of competitive local-exchange carriers (CLECs), were major catalysts of change. As a result, the domestic competitive landscape continues to accelerate while global competition is increasing due to World Trade Organization (WTO), European Union (E.U.), and other country directives.

Technology and market forces have also unleashed a whirlwind of change, including the recent dramatic growth of the Internet, the time division multiplexing (TDM)–to-packet network evolution, and an increased demand for broadband services. To meet these demands, a major transition from circuit-switched technology to packet-switched technology is occurring.

Packetized voice brings the advantages of lower capital and operational expense. It also has the ability to offer enhanced services cost-effectively and rapidly. Furthermore, it offers the efficient transport of voice and data on a common network infrastructure while providing increased networking flexibility and forward migration capability.

This transition is being reflected in market-analyst reports. In an August 2000 study by the Yankee Group, analysts concluded that service providers must provide a suite of integrated offerings including voice, data, and other broadband services while simplifying customer interaction and support. According to this study, network solutions should also provide the means for provisioning services and billing in a convenient manner.

Also driving the transition from circuit-switched technology to packet technology is the explosive growth in demand for high-speed data, both in the backbone of the network and in the local access loop. But, as reported in a front-page article in the *Wall Street Journal* (May 11, 2001), some 97 percent of the optical capacity of the core network is not being used today. This is due, in large part, to a lack of multiservice local access exchanges that could fill this gap by enabling a variety of advanced broadband offerings to move up stream.

As a result, both transport and converged switching capabilities are in high demand in the local-exchange and access network in order to adequately justify the high-bandwidth pipeline in the core network.

Demographics and cost-of-living economics are also contributing to the distribution of broadband services, as customers continue to move to the suburbs or rural areas. These customers are increasingly demanding advanced services. These demands are placing a heavy burden on the COs as local service providers attempt to find cost-effective solutions. According to one industry luminary:

"The first mile remains an important bottleneck. Fiber is still too expensive compared with copper, except where very high-speed applications are present… The core of the network has become a hybrid of wavelength (optical) switching and massive IP routing."

—Dr. John McQuillan
NGN Ventures, 2001

Historical Perspectives

From 1900 to the 1970s, beyond the provisioning of basic service, there were very few improvements in customer applications. Later on, Centrex features and custom calling services, or custom local-area signaling services (CLASS), provided by the local CO and new, colorful desktop/wall-mount customer-premises equipment (CPE) were the focus. On the business side during this period, upgrades included a wide range of private branch exchange (PBX) and key system features. Again, these were primarily focused on the provision of voice-grade services. Early forays into data services, such as the integrated services digital network (ISDN), were very expensive to implement, and high monthly subscription rates—coupled with a real lack of applications—kept widespread adoption at bay. In the final analysis, dial tone has not changed in more than 100 years.

At the same time, internal telephone company plant equipment, switches, and transmission facilities were being built at a hurried pace to match the rapid adoption of basic telephone service.

At the heart of the public switching network, the legacy Class-5 office handles local call switching and customer application features. The Class-5 office is the critical, initial point of access to the network for every subscriber line connected to the local exchange.

As the port of entry for all users of the public switched telecommunications highway and the originating source for dial-tone and telephony services, the end office is in a unique position. It influences all calls channeled through the various switching layers of the national hierarchy.

It is at the local CO, recently coined the "service point of presence" (SPOP), that cost-efficiency can be maximized and service provisioning and maintenance services can be optimized, with the help of intelligent softswitch functionality and packet connectivity layers.

Local Switching Adoption Cycles

Before exploring the critical, new role that Class-5 end-offices play in the emerging network, a brief overview of the five local switching adoption cycles that have occurred since the late 1800s is in order. While each cycle had a definite start date, in most cases, succeeding generations of switches overlap prior ones by a number of years, existing side by side at different locations. Each major cycle change has been a once-in-a-lifetime event (see *Figure 1*).

Cycle 1

In 1891, an undertaker by the name of Almon B. Strowger developed a dial-pulse–driven, electromechanical (or step-by-step) switch in response to his personal desire to thwart rivals from learning about his new business leads through operator-assisted calls.

Although he was not a telephone engineer, he reasoned that if users could dial calls directly, it would eliminate the need for an operator, and thus save his sagging funeral business from encroaching competition. He believed that the operator was diverting calls to his hometown competitor.

The "stepper," as it was called, was rugged and reliable. In fact, it was so robust that it remained in service as the principal network switching system for more than 75 years.

It is interesting to note that, from the beginning, innovation and early adoption of new switching and other technologies did not come from large, established service providers, but from smaller, independent phone companies and entrepreneurs.

Cycle 2

1938 witnessed the introduction of a common control electromechanical switch, known as the No.5 Crossbar, first installed in Brooklyn, New York. The crossbar switch greatly increased processing speed and switching volumes but retained overall system "intelligence" in the hardware.

FIGURE 1

The Five Adoption Cycles of Switching

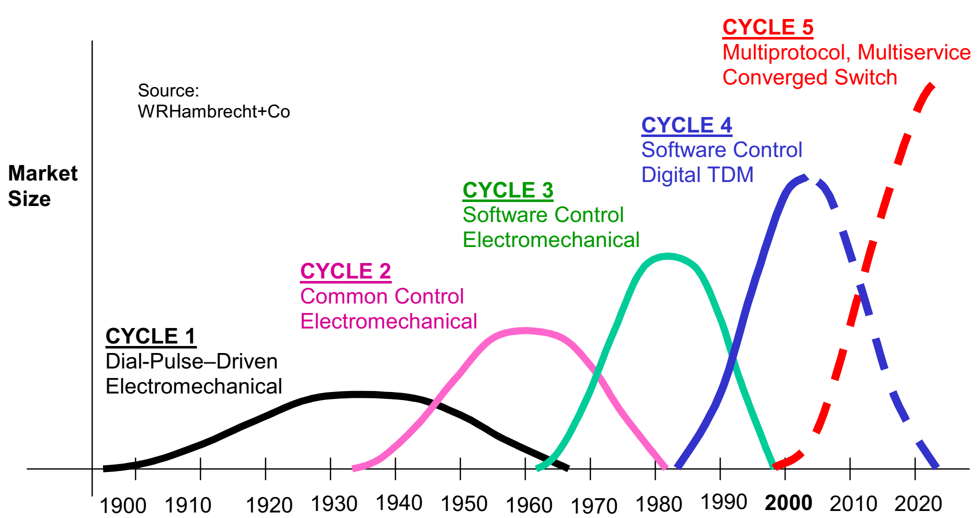

Cycle 3

By 1965 computing technology had advanced to the point where the Bell System could deploy the first software-controlled electronic switching system, known as the No.1 ESS, initially in two New Jersey locations: Succasunna and Trenton. This system became the principal switching platform used in COs for some 35 years.

Cycle 4

The late 1960s and early 1970s saw the dawn of the Internet and the concurrent rise in digital data technologies. This led to the introduction of a new breed of software-controlled, TDM (64 kbps timeslot) switches. Known as the No.5 ESS, this switch is still considered to be the workhorse of the CO today.

As Internet traffic volumes exploded in the 1990s, service providers searched for alternative switching solutions to help offload the heavy Internet traffic volume while also paving the way for advanced data and broadband services and other sources of new revenue.

By the end of the 20th century, the voice market represented some 80 percent of total revenue, but voice-service growth was slowing. At the same time, demand for data services was increasing, but legacy switches were not designed to process digital data beyond 64 kbps, and revenues from emerging data services were not filling the gap. The stage was set for an entirely new breed of converged voice and data local-exchange switch, with scaleable capacity, to offer a host of new services.

Cycle 5

Today. Next-generation (fifth adoption cycle) local switches entering the network in the new millennium must be designed to handle the transition from a circuit-switched (legacy voice, TDM, and public switched telephone network [PSTN] services) world to a packet-switched network environment. This new "converged" next-generation switch must be able to handle the following:

- Data services, up to the gigabit Ethernet broadband service level
- Multiprotocol services, including the following:
  - Internet protocol (IP)
  - Asynchronous transfer mode (ATM)
  - Synchronous optical network (SONET)
  - Base revenues derived from legacy TDM voice

Three Transitional Influences

According to Bear Stearns analysts in *Saving Telecommunications: The Next-Generation Access and Services Evolution* (March 30, 2001), "Today, the three primary transitionary influences are wreaking havoc on the telecommunications industry. These three influences can be summarized as follows:

1. We are transitioning to a period where the focus must be on *access modernization* versus core network modernization,
2. We are transitioning towards *next-generation networking* (in all segments of the network) from old-generation networking, primarily outside of the core, which has already moved so destructively down the modernization path, and
3. We have the most fundamental shift of all taking place today (*the move to data services from voice services* and thus to broadband networking from narrowband networking."

One of the key operative terms in this research report is *"access modernization"* (translated as the Class-5 local-exchange upgrades and physical copper lines). We are already transitioning toward *"next-generation networking"* in all segments of the network—away from old-generation networking primarily outside of the core, which has already moved destructively down the modernization path. Furthermore, the most fundamental shift of all is taking place today—*"the move to data services from voice services"* and, thus, to broadband networking from narrowband networking.

This next-generation solution must, of necessity, involve a converged voice/data solution while combining legacy and advanced services in a single package that is both backward and forward compatible.

The new integrated communications providers (ICPs) are experiencing similar challenges due to the number of subscribers needed to ramp up their business. At the same time, there are strong forces demanding early profitability. The average number of ICP digital subscriber line (DSL) lines in any given co-location is low, and voice services typically must be profitable while ramping from 0 to 20 percent market-share in targeted markets.

In the incumbent–CO as well as the competitive-carrier situation, there is a compelling need for a system that provides low-to-medium line-count cost-effectiveness. For ICPs, this requires a high degree of scalability to handle the extremes from a few subscribers up to, and including, high-volume growth. In this scenario, a scaleable approach addresses the low–line-count cases.

Furthermore, there is a good opportunity in the present mode of operation to increase efficiency and simplify operations by integrating once disparate systems into a single, converged system.

The Challenges of Today's Infrastructure

In summary, traditional networks over the years were designed and built based on the need to provide reliable voice services. As demand for data applications has exploded, there has been an aggressive buildout of data networks to overlay or co-exist—as in the case of DSL, DSL access multiplexers (DSLAMs), and edge ATM—with the voice network. This has resulted in the following:

- Multiple systems that must interoperate; requiring extensive space, power, human resources, and time
- Disparate network-management and billing systems
- Switching systems optimized for very large deployments that do not scale well for smaller markets

The existing network infrastructure has evolved from a voice-only network to a high–data-traffic network in recent years. This rapid expansion of data services has resulted in a separate (overlay) network on top of the existing voice architecture.

While this multibox system has supplied the means for providing new services, it has proven itself to be an inefficient solution—since at least two networks are necessary for total service provision to a single end user. Today, this conglomerate solution is in danger of collapsing under its own weight.

Legacy Class-5 TDM Switches Must Be Replaced

The case for replacing existing Class-5 switches can be supported from many vantage points. First, as has been discussed, they are not data friendly, nor are they broadband capable. Second, the cost involved to upgrade the current generation is prohibitively high. Even if this could be done, today's CO switches are too large, their software pace is too slow and operational, and maintenance expenses are too high.

All things considered, there is an even more urgent reason to adopt a total replacement strategy: Several manufacturers, including Lucent Technologies, Nortel Networks, Ericsson, and NEC have announced that they will discontinue supporting their small CO products.

Since 85 percent of U.S. COs serve fewer than 4,000 lines, this means that most of the U.S. telephone network is operating on obsolete systems. Internationally, 90 percent of local exchanges serve 3,500 lines or less. The end result is that almost all of the world's population is connected to, and uses, an obsolete PSTN.

Criteria for a Next-Generation Converged Local Exchange

Local service providers must find a cost-effective broadband solution for suburban and rural COs—one that will continue to support the legacy services that provide the bulk of their revenue while using minimal CO space. Furthermore, this new switch must be capable of providing robust, always-available connectivity in emergencies and be highly responsive to the real-time service provisioning required to meet end-user demands. Beyond these basics, such a solution must minimize the need for up-front capital, installation, configuration, and ongoing maintenance costs. Putting one "box" in the place of many sounds like a tall order, especially at a time when service providers are concerned about losing their TDM revenue.

The answer lies in the prospect of finding a packet-capable voice and data-networking device that can furnish the necessities at an easily affordable price. The replacement issue goes directly to the bottom line. Independent operating carrier (IOC)/ICP customers can no longer afford to make use of current switching solutions and expect to maintain profitability and continued revenue growth for broadband services.

The Need for Softswitch Technology in Fringe Areas

There is much discussion today over whether or not softswitch intelligence should reside in the core or at the edge of the network. The term "softswitch" refers to software that provides call control, routing instructions, and features in a next-generation voice network. A softswitch (also known as a call agent or media gateway controller) connects to a media gateway and directs calls through a packet-based network. Proponents of a centralized softswitch approach are attempting to offer solutions that fill the core with converged switching solutions but do not address key local office and access issues.

With more and more people moving from cities to suburban or rural areas, and others spending more work hours telecommuting from home or using the Internet, there is an increased demand for broadband services in fringe areas. The problem is to remove the existing bottleneck and fill high-bandwidth systems in the core network. It is interesting to note that these are the very same geographic areas currently served by the obsolete small exchange.

A growing number of service providers now believe that these new, so-called core-centric architectures are flawed. They embrace the notion of a "smart edge" with a "dumb core." Supporters of this concept argue that there is a need to push softswitch intelligence down to the SPOP to avoid reliability issues and to ensure end-office survivability in emergencies. In moving intelligence away from the edge, local COs will lack the ability to stand alone, providing mandatory services such as E-911, call processing, and service provisioning.

Furthermore, there is a mandate in all areas for local law-enforcement authorities to have the ability to impose lawful intercepts under the Communications Act for Law Enforcement Agencies (CALEA) when required to fight crime, including illegal alcohol and drug trafficking. With softswitch intelligence in the core, complete CALEA and 911 requirements cannot be met. A new, 21st-century end-office switching criteria can only be provided using an entirely new breed of converged local exchange.

Big City Solutions Do Not Fit the Small CO Model

Some have proposed taking switches designed for large metropolitan-area markets and adapting them for service in small-town COs. On the surface, this idea may appear to have merit, but a closer look leads to an opposite conclusion.

Solutions that repackage the larger switch into a smaller, line-size switch with added packet services may appear, on the surface, to make sense. But because end-users served by small COs still want the same services as those available in the big cities, these solutions involve multiple-box, multi-rack solutions.

This dilemma is compounded because businesses want multi-line, low-cost voice and data access, and both residence and business customers want high-speed Internet access. Both of these groups require services in emergencies.

The Answer: "CO in a Box" or "Softswitch in a Hard Shell"

The ideal solution is to provide CO services in a single box—a small, scalable multiservice, multiprotocol system that stands alone as a basic switch, interworks with the packet and softswitch core, and provides legacy and advanced services. Until recently, attempts at delivering this type of product have only produced partial solutions.

Gluon Networks, Inc. has designed and developed a single, converged local-exchange solution that meets the criteria outlined in this paper. This solution, called the converged local exchange (CLX), offers the value proposition of increased revenue at lower cost—revenue from voice and data with a single infrastructure at 1/6th the capital costs, only 1/20th of the footprint now required, and only 1/10th of the average time-to-revenue.

This product is a small, flexible carrier-class voice and data switch that combines the functionality of a Class-5 voice switch with physical access, DSLAM, an ATM switch, SONET, and packet transport.

Asia and Australia Telecommunications Industry Overview

Paul Budde
Telecommunications Consultant
Paul Budde Communication

1. Asian Market Overview

1.1. Asian Telecommunications Champion of the World

Whether you are talking about mobile growth, Internet expansion, submarine cable projects, or fixed phone lines, Asia now leads the world in every segment of the telecommunications market.

In most of these markets, growth in Asia is twice, sometimes three times, that of Europe or North America. The developments are widespread:

- South Korea is leading the world in broadband growth.
- Japan is the world leader in mobile applications.
- Asia as a whole is leading global Internet growth.
- China leads the world in sheer subscriber number growth.
- Hong Kong is positioning itself as the international telecom hub.
- The Philippines and Taiwan are championing short message services (SMSs).

There is no obvious telecommunications recession in Asia, and, with the ongoing exception of Indonesia, growth is spread throughout the region—perhaps not always equally, but far more widespread than, for example, in Latin America and Eastern Europe. India appears to have at last turned the bureaucratic corner and is now better positioned to kick off on a potentially significant growth spurt. In virtually all segments of the telecommunications industry, pent-up demand amongst its middle class is larger than the total telecommunications market of many European countries. India will soon follow China in an unprecedented growth in telecommunications.

And, as we have indicated in previous years, Asia is still in the early stages of its information and communications boom. Asian people are amongst the most enthusiastic users of communications and information in the world. Apart from the obvious commercial opportunities, other sectors, such as education and community services, are seeing incredible growth.

In an industry dominated by Western cultures and values, global statistical overviews are always based on the Western principles, ignoring the differences inherent in the societal structures of the Asian people. In Asia, it is not essential for every individual in a family or community to own his or her own mobile phone, Internet connection, and television. Often, one single device services from 10 to 50 people. Based on this principle, well over 50% of all adults in Asia have access to mobile telephony, Internet, television, etc. Again, this is higher than in several of the Western countries. Furthermore, current developments are producing an ongoing increase in this high penetration figure. Yet, the official statistics continue to talk about teledensity (telephones per 100 people) levels of 10% to 20%.

Western-based companies still dominate the technology supply lines, and most of these fail to adequately address community-based needs. In general terms, they continue to take an individual approach. However, companies that are capable of changing their business models to incorporate products and services that allow sharing of networks, devices, and services will do very well. Although Asia also may develop over time into a more individual-oriented society, this could take as many as 25 years.

On the other side, after the dramatic terrorist activities of 2001, there is also the possibility, of course, that the West, in its own search for social solutions, will adopt some of the more community-oriented ways of living, especially if such alternatives are supported by the technological product developments coming out of Asia.

Nevertheless, there will always be a distinct cultural difference between the West and the East.

2. Australian Market Overview

2.1. The Collapse of Competition

Developments during the first half of 2001 can only be described as disastrous for competition in telecommunications in Australia. There is a more than 50% chance that the country will end up with a duopoly in this industry, with Cable & Wireless Optus being the only certain survivor to Telstra.

Overall growth in the industry has dwindled from 14% in 2000 to –0.1% in 2001 but will slightly recover in 2002 to 6% (see *Table 1*). It will not be until 2003/2005 that broadband will have caught up sufficiently to lead the market into new growth. As there will be significant pent-up demand for both broadband access and broadband services, spectacular growth is anticipated in the second half of this decade, resulting in a $90 billion[1] market by 2010.

According to the findings in our Telecommunications Industry Report 2001/2002[2], the new telco industry will not grow to the initially expected $5 billion-plus market this year; instead, it will contract by 5% to a level of $3.5 billion.

After the spectacular collapse of One.Tel, another champion of competition bit the dust. WorldxChange single-handedly changed the nature of competition in the long-distance telephone market in Australia. It guided the industry to the level of call charges that we currently enjoy. Hutchison sold its resale business and thus also dropped from the list of major telcos (annual revenues of $100 million).

At the same time, companies such as PowerTel and UECom failed to live up to their promise and never created a real impact on the market in Australia. WorldCom, another CBD operator, has been dormant for nearly 18 months and is only now emerging from its long sleep. This has also resulted, finally, in a revival of OzEmail. This promising Internet service provider (ISP) lost its number-one position to Telstra and is now a very distant number-two player in the market. It has also lost ground in the emerging broadband market.

While there are now more than 60 carrier service providers, the top six account for 95% of the total market. All the regional operators together have less than a 0.5% market-share.

3. Asian Fixed Network Market

Projected demand for telecommunications services in Asia indicates that at least U.S.$1 trillion will be invested in new infrastructure over the current decade. Regional governments have recognized, and consistently placed an emphasis on, the role of telecommunications in national development. Even before the Asian economic crisis, they were acknowledging the need to encourage private-sector investment to help meet the shortfall in investment capital. This involvement by the private sector is now even more vital, with an increasing interest in attracting foreign capital becoming obvious.

Although infrastructure developments slowed with the economic crisis of the late 1990s, demand for wholesale services has risen. Driven in the short term by voice services, but in the longer term by data services (such as the Internet), demand for wholesale services is expected to treble by 2002. Planned fiber network build-outs in the region should boost bandwidth 20-fold. The year 2003 will see close to 90% of all new connections made through optical fiber-to-the-home (FTTH) or fiber-to-the-curb (FTTC). Already by the middle of this decade, it is estimated that Asia will lead the FTTH market, as the operators of legacy networks in other parts of the world will try to protect their vested interests in copper-based networks.

Telecommunications emerged largely "crisis proof" during the dark days of the Asian financial meltdown due to strong pent-up demand for wireline and wireless services. Growth in main telephone lines during this period slowed only slightly and remained at 10%, which was twice the world average. Telcos in some of Asia's developing economies (China, India, Indonesia, Malaysia, the Philippines, Thailand, and Vietnam) serviced 90% of total demand for telephones lines (150 million), an indicator of Asia's appetite for telecommunications.

The decade of the 1990s saw a three-fold increase in the number of telephone users. Asia now accounts for a third of the world's telephone users.

One problem for the established fixed network operators in Asia has been voice over Internet protocol (VoIP) telephony. Without any real investment in infrastructure, an interesting mix of licensed and unlicensed service providers across the region are beginning to offer heavily discounted national and international VoIP–based telephony services. The entry of these operators into the market poses a serious threat to the revenues of the mainstream telcos. Possibly more significantly, it has raised fundamental questions about the adequacy of national policies and regulatory regimes governing telecommunications.

4. Australian Fixed Network Market

The country is currently peddling backwards. Telstra, through yet another appeal, wants to increase the wholesale rates. This is diametrically opposed to developments that are taking place overseas. They are also using a loophole in the legislation to refuse termination on their network from certain competitors; Telstra was the company that lobbied the hardest for the current termination rules in the first place! Asymmetric digital subscriber line (ADSL) and local-call wholesale rates are totally unsustainable competitively, and the government has been merely paying lip service to the problem and promising changes to the regime to stop Telstra from delaying competition, year after year after year.

TABLE 1

Revenue Growth in the Australian Telecommunications Market, 2000–2002

Year	Fixed voice	Mobile	Data	Equipment	Total
2000	10.2%	12.4%	24.0%	11.7%	14.2%
2001 (e)	–2.1%	7.9%	18.0%	–12.5%	–0.1%
2002 (e)	–3.8%	–5.4%	27.2%	10.2%	5.7%

Source: Telecommunications Industry Report 2001/2002; Paul Budde Communication

Even if their latest appeal is not upheld, it could delay competition for another two to three years, and by that time, we won't have to worry any more because the market will have been re-monopolized.

4.1. Flawed Government Policies

While it would be easy to blame all of this on the Australian regulatory environment, it has certainly played a key role in the demise of competition in these cases. While some of the business plans from the aforementioned players have not been particularly sound, all of them had built into their plans better wholesale prices and access to customers—perhaps not immediately after deregulation in 1997, but at least one or two years later. As this didn't happen, they could no longer sustain the unprofitable wholesale margins that they received, and this eventually led to the downfall of the majority of them.

I was called a pessimist in 1997 when I indicated that Telstra could delay the introduction of competition by two years. Now, five years later, none of the key issues (interconnect, access to the local loop, and number portability) has been implemented. The technical regulator ACA has gone well beyond the call of duty to secure the introduction of number portability. Yet without a strong regulatory regime, it will not be able, even for those companies with the best business plans, to successfully operate in Australia.

5. Asian Internet Market

Internet growth in Asia continues to be led by the developed economies of the region: Australia, Japan, New Zealand, Hong Kong, Korea, Singapore, and Taiwan. By contrast, in the developing economies, Internet usage is only inching forward due to the high cost of access, poor telecommunications infrastructure, and the sluggish pace of deregulation. Nevertheless, even in the developing economies of Asia, promising economic returns generated by the Internet is compelling both government and private enterprise toward modernization. There is simply no other choice.

Asia made important strides in satisfying its huge appetite for the Internet in 2000. The world's third largest regional Internet market after the United States and Europe saw significant increases in the number of Internet users and e-commerce revenues during the year. More importantly, Asia continued to move rapidly ahead in high-speed Internet access, increasingly the preferred mode worldwide. South Korea had more broadband Internet subscribers than the United States, a country with five times its population.

6. Australian Internet Market

With such high stakes, carriers will try to dominate the Internet market. Tto survive, let alone flourish, ISPs will have to join forces within the Internet market as well as within the total telecommunications services market. The ISP market will grow from $1.2 billion in 2000 to $3.5 billion in 2002. The bigger companies are, in general, growing faster than the smaller ones. However the smaller ones are more profitable. There has been a marked improvement in operational margins. However, for the larger companies, this is not enough to cover their sales and marketing costs. Those ISPs that have an early start in ADSL will pick the results in 2002/2003 (iPrimus, Pacific Internet).

Overall developments in this market will eventually lead to the following:

- Half a dozen large ISPs
- Between 20 and 30 business market ISPs looking after market segments such as travel, education, games, government services, and financial advice services
- Perhaps as many as 50 to 100 ISPs servicing regional markets
- A large number of franchisees, agents, etc. (hundreds)

TABLE 2

Top Asian ISPs and Subscribers

ISP	Country	2000
iMode	Japan	8,289,000
Nifty-Serve	Japan	3,700,000
Cholian	South Korea	3,134,000
Unitel	South Korea	3,111,000
Hitel	South Korea	3,030,000
Biglobe	Japan	2,990,000
EZWeb	Japan	2,691,000
n. Top	South Korea	2,200,000
Netsgo	South Korea	1,745,000
J-Sky	Japan	1,744,000
Total		**32,633,000**

Source: Telecommunications and Information Highways in Asia Report, Paul Budde Communication 2001

Note: Throughout Asia there is also a strong focus on providing public access to the Internet. Cybercafes are mushrooming in the major cities and even pushing out into the rural areas. In India, for example, it is estimated that a massive 60% of Internet users go on-line at the public access points.

Other ISPs might find niches in certain nonprofit, educational, and hobby markets. The vast majority of the current ISPs will disappear or will become agents and franchisees for the larger companies. Despite various reports on imminent shakeouts of the markets, the ISPs have so far proved very resilient; however, a steady decline in numbers began to occur over the 2000–2001 period.

With the discussion broadening on how to best implement the Internet, the focus is shifting from simply providing connection to offering value-added services. This, in its turn, is going to separate the men from the boys.

The key to this will be streaming media. We expect this market to take off commercially around 2003. By that time, we not only expect the technologies needed for this to be robust enough to operate within large-scale broadband networks, but, perhaps even more importantly, we expect that more bandwidth-sucking applications (BSAs) will become available. As soon as customers become accustomed to high-speed, always-on Internet services, they will be ready for new services (e-entertainment, home video, pictures, teleseminars, etc.). A problem for many of the dot-com content providers is their inability to distribute video-based content.

7. Asian Mobile Market

The mobile industry in Asia remains one of the fastest-growing markets in the world, with almost 234 million mobile phones by year-end 2000 and revenues of more than U.S.$100 billion in mid-2000, compared with U.S.$10 million in 1994. Revenues will be close to U.S.$150 billion in 2003. Investment between 1999 and 2004 is expected to be in the range of U.S.$300–600 billion. According to the International Telecommunications Union (ITU), the region is poised to become the world's mobile powerhouse. It predicts that by 2010, more than 50% of all mobile phone users in the world will be in the Asia-Pacific region, up from a 35% share in 2000.

Despite the economic downturn in the region, Asia had annual wireless growth rates of 50% in 1999 and 2000. Seriously affected by the economic crisis, Indonesia's mobile market barely grew by 16% in 1998, but soared more than 100% in 1999 and more than 62% in 2000.

Wireless technology holds much promise for promoting communication access in the region. Cambodia was the first country in the world to boast more mobile than fixed subscribers by 1993, and there are now more mobile than fixed phone subscribers in a number of other Asian countries, with several more very close to crossing the line.

The region is home to the two largest mobile companies in the world: China Mobile and Japan's NTT DoCoMo. During 2000, China moved ahead of Japan as the world's second biggest mobile market. At the current rate of growth (more than 65%), China will soon overtake the United States as the biggest market in the world. In markets such as Indonesia and the Philippines, growth has been mainly driven by prepaid services. In Japan and Korea, new Internet services have been the catalyst, whereas in Singapore and China, increased competition has been the prime mover.

7.1. Asian Mobile Data Market

Wireless data subscribers in Asia accounted for 39% of the world market at year-end 2000. Japan is expected to account for a quarter of the world total by 2003. The introduction of

TABLE 3

Internet Market Revenues, 2000–2002

Year	Total revenue ($ million)
2000	1,200
2001(e)	2,500
2002(e)	3,500

Source: Telecommunications Industry Report 2001/2002, Paul Budde Communication

TABLE 4

Fixed to Mobile Telephone Substitution (Selected Countries)

Country	Date Mobile Passed Fixed	Subscribers (millions)	
		Mobile	Fixed
South Korea	08/1999	26.7	21.9
Hong Kong SAR	11/1999	5.3	3.9
Taiwan	02/2000	17.6	13.0
Japan	02/2000	58.0	55.5
Singapore	07/2000	2.4	1.8
Philippines	07/2000	6.0	3.0

Source: Telecommunications and Information Highways in Asia Report, Paul Budde Communication, 2001

TABLE 5

Top 10 Asia-Pacific Mobile Operators by Subscribers

Operator	Country	Mobile Revenue (U.S.$ million)
China Mobile	China	8,700
NTT DoCoMo	Japan	35,000
SK Telecom	South Korea	3,800
J-Phone Group	Japan	6,600
China Unicom	China	1,000
DDI Group	Japan	7,700
KT Freetel	South Korea	1,300
Telstra	Australia	1,700
Chunghwa Telecom	Taiwan–China	1,400
IDO	Japan	n/a

Source: Telecommunications and Information Highways in Asia Report, Paul Budde Communication, 2001

third generation (3G)-related services in Japan during 2001 will be the principal driver behind the growth in the number of wireless data subscribers. IDC forecasts 142 million wireless Internet users in the region (excluding Japan) by 2004.

With the possible exception of Japan and South Korea, the initial rollout for 3G services in Asia will most likely follow the pattern set forth by wireless application protocol (WAP), with expectations far exceeding reality. User acceptance level is likely to be relatively low due to the high price of 3G handsets and the initial speed for 3G-related services failing to meet expectations. Another important issue is how many useful applications will become available for 3G when it is launched. A wide range of value-added applications is critical to increase wireless data usage.

8. Australian Mobile Market

8.1. Another Two-Horse Race

The collapse of One.Tel was nothing less than spectacular, but it has to be noted that this had more to do with the involvement of Packer and Murdoch than with the life and death of the telco itself. True, in the end the directors are the ones that have to bear the responsibility and, as we indicated 18 months ago, they went off the rails when they changed their business model from being a virtual telco operator to being a carrier.

Different business models are needed for this, and, as was proved in the long run, they could not be rolled into one. Having said this, there was at one stage plenty of room to change the business model back to its original form, and One.Tel could have been saved. One.Tel has been critical to progress of competition in Australia, and after its demise we have seen a dramatic downturn of competition in the mobile market—with Telstra increasing prepaid prices and Vodafone severely curtailing its rollout of services to customers in the lower end of the market.

8.2. Mobile Growth Comes Again to a Grinding Halt

Once again growth in the market has come to a stand still. By mid-2001, the triopoly called it quits—no more aggressive drive for new customers, reduction in the acquisition of new customers at the bottom end of the market, and a consolidation of their average revenue per user (ARPUs). Aggressive competition during 1999 and 2000 saw the penetration of mobile phones increase from 40% to 60%. The cellular companies (celcos) have been heavily involved, particularly in 2000, in price competition that saw their mobile revenues per user dropping quite dramatically.

The prepaid market, bringing in new customers, has been the main driver of growth; however, many customers are youngsters and teenagers with a very low spend per month. Prepaid now accounts for the majority of new subscribers to the network. Their annual spend is often less than $250. This, of course, is a drain on the celcos, as their organizations are not equipped to deal with such small customers. Their overheads are far too high to operate profitably in such markets. This is yet another reason why celcos should look more closely at opening up their networks to wholesale customers, who would be much better positioned to service specific markets. The development of mobile data is another reason to use the virtual mobile networks operators (VMNOs) model, as specialized companies are far better equipped to develop this market. However, with the exception of C&W Optus, none of the others are keen to allow these virtual mobile networks operators onto their network.

As for 2002, we predict a negative overall growth in the mobile market as the products become commoditized. This will continue for another two or three years until the market begins to stabilize around 2005. It will be interesting to see if new business models will have been developed (allowed) by that time; only these new models will bring back the revenue growth in this market. It looks like industry restructuring cannot be done on a "voluntary" basis; in Europe, it is currently the financial market that is forcing such changes through in order to get the operators to reduce their huge debts.

The future of mobile will not be in mobile data; this will remain a niche market. Low-level data services—such as SMS, competitions, ticketing, and transfers—will provide some revenue relief; however, these services will mainly be

TABLE 6

Average Revenue per Mobile Customer, 1993–2005

Year	ARPU p.a.
1993	$1500
1995	$1000
1997	$800
1999	$780
2001 (e)	$630
2003 (e)	$610
2005 (e)	$680

Source: Telecommunications Industry Report 2001/2002, Paul Budde Communication
Note: Excluding handsets and fixed to mobile revenues

offered as promotional services to attract new customers and/or to stop churn. High SMS traffic also requires a better mobile infrastructure to cope with the increase in demand. This, of course, adds extra costs to what often prove to be marginal services. Nevertheless, having said this, over time these services could make up between 15% and 20% of all mobile revenue.

9. Asian Broadband Market

Asia looks to broadband to efficiently deliver the next generation of Internet services, such as interactive television and streaming media. However, broadband has so far only managed to make an impact in Asia's developed economies. Everywhere else in the region, dial-up narrowband access predominates, at best achieving 56 kilobits per second (kbps).

Broadband subscribers in Asia are forecast to increase from less than 500,000 in 1999 to about 35 million by 2003.

The potential of broadband as the premier Internet delivery technology is being held in check by its high cost. However, in countries with strong broadband demand, such as South Korea, the issue is not so much cost as the shortage of suitable high-speed local language content.

With an estimated 2 million broadband subscribers by the end of 2000, and a growing rate of 25% per month, South Korea is the world's leading broadband country. The high density living patterns found in South Korea have made broadband more cost-effective than in other markets. Forecasts are for 15 million digital subscriber line (DSL) subscribers and 2.4 million cable-modem users by 2004.

Although Hong Kong has one of the world's most advanced and sophisticated telecommunications infrastructure, with 100% of the territory and 90% of all homes covered by broadband connections by the close of 2000, growth in broadband subscriptions has been relatively slow. By February 2001, there were only 417,000 broadband customers from a population of 6.7 million, compared with 2.2 million dial-up subscribers.

In Singapore, almost 100% of the island's homes and businesses are passed by broadband infrastructure. There were an estimated 275,000 broadband users in February 2001, up from 100,000 users at year-end 1999.

China had an estimated 175,000 broadband subscribers at the end of 2000, as opposed to 880,000 narrowband subscribers. Broadband growth in China is projected to surpass that of narrowband by 2002 and should reach close to a million subscribers by 2004.

With 92,000 DSL subscribers and 72,000 cable-modem users in Taiwan in 2000, the country is set to have 1.5 million DSL subscribers and 410,000 cable-modem subscribers in 2004.

Somewhat surprisingly, broadband is still relatively in its infancy in Japan, with little more than half a million people using broadband access through cable or ADSL in early 2001. But the government's e-Japan Priority Policy Program has set the goal of having five million broadband subscriptions by 2002 and 30 million broadband households by 2005.

10. Australian Broadband Market

10.1. The Broadband Crusade

During 2000, it became abundantly clear that broadband was not growing as fast in Australia as it was in other parts of the world. So I decided that 2001 would be my Year of the Broadband Crusade. The main reason for the crusade was the announcement in late 2000 that Telstra was predicting a meagre 625,000 ADSL users by 2005 (one million if you include cable-modem users).

I am clearly disappointed with Telstra's one million-plus forecast, especially in view of the fact that the company actually is rolling out broadband networks.

I see broadband as a crucial factor in the development of our country, from both a community and an economic perspective. In comparison to other countries, we are currently running 18 months behind on the information superhighway, and if Telstra's plan succeeds, this gap will increase to three years by 2005.

Broadband is more of a concept than a defined technology. It provides convenient and affordable access to information, entertainment, and communication (ADSL, cable modems,

TABLE 7

South Korea Broadband Users, 2001

Access Method	Subscribers
ADSL	3,469,600
Cable Modem	1,944,100
Satellite	9,000
Total	**5,422,700**

Source: Telecommunications and Information Highways in Asia Report, Paul Budde Communication, 2001

integrated services digital network (ISDN), and even an always-on flat rate 56K modem service could be seen by some customers as a broadband service). Once video-based content becomes more widely available, customers will automatically have higher broadband expectations, and so higher-speed services will increasingly be linked to their concept of broadband (45 megabits per second [Mbps] by 2010). Broadband empowers individuals to choose the services they want, rather than being restricted to what the media and telco barons are prepared to dole out to them. It is also the infrastructure key to the concept of a knowledge-based society and an e-economy.

So versatile is this technology that it can be applied in many ways, in accordance with user demand. For example, in Korea, the country with the highest broadband penetration in the world (57% penetration amongst Internet households), the market is driven by tele-education. In the United States, teleworking is operating as a key driver, while always-on, high-speed Internet access is the favourite application in Europe. Gambling and porn are also very strong market drivers.

The impetus for broadband can occur in three ways:

- A national vision translated into government policies (the governments of the Netherlands, Sweden, and France are actively involved in the rollout of broadband), possibly leading to the concept of LoopCo.
- Bold telcos and cable operators leading the charge (Spain, Germany, Korea)
- Competition between delivery platforms such as telecoms, cable television, and digital television (United States, United Kingdom, Hong Kong)

Australia does not have any of the above drivers, hence my crusade, in which I have endeavored to open up the subject by using the media to generate public awareness and debate. The ultimate defeat of the current generation of telcos would be if a future government has to nationalize the telecoms infrastructure (LoopCo) in order to secure a first-call infrastructure for our e-society and e-economy.

"LoopCo will have to be separated from the rest of the telco business."

11. The Crash of the Equipment Market

Disaster struck in 2001, when this segment of the industry contracted by more than a quarter. Total Australian equipment revenue (telco and datacoms) will fall to $8 billion instead of growing to $11 billion as had been forecasted. While the market will recover, major structural damage will have been done to the vendor industry. While most of the big six global vendors will survive, a major restructuring will be needed, during which, among other things, the issues subsequently listed will need to be addressed.

The key elements of the crash were as follows:

- Hyped up technology expectations, especially in the mobile market and long-haul access market
- Collapse of a totally unsustainable dot-com market in the United States, fuelled by a greedy investment market
- Lackluster development in local broadband access
- Promotion of big iron technologies, such as long-haul telecommunications infrastructure instead of lean, mean access technologies for residential and small-to-medium enterprise (SME) customers
- Holding on to outmoded business models linked to "old" telco models

It is sad to see that the immediate reaction to the crisis focused on accountancy solutions aimed at protecting the current pie. Only a few are trying to work through the emergency through growth (Nokia and Marconi). Lucent Technologies has so far been the only one that has announced a major restructuring. Again, others, such as Siemens and Philips, have been less affected, because their communications divisions are supported by other activities.

As in the operator industry, massive restructuring along horizontal lines is needed to turn the industry around.

12. New Generation of Telcos

As the first and second generation of telcos are faltering, a new breed of competitors is entering the market. Some have re-engineered themselves; others are first entrants. They are as follows:

- Broadband telcos
- Virtual operators
- Regional operators
- E-business service providers, application service providers (ASPs), service bureaus, and data-center operators

Their business models are not formulated along the same lines as those used by the "old" telcos but on specialization into key market segments, such as regional markets, broadband, wholesale infrastructure, etc.

Furthermore, voice telecommunications constitutes only a minor part of their business models; they are concentrating on data services, content delivery, and multi-access technologies.

It looks as though the old competition model is dead but not yet buried, and it will take a few years before the new generation of telcos establish themselves. This will also depend on regulatory issues, most importantly the local-loop unbundling issue.

Already re-engineered companies include Premium Rate service providers (audiotext services) and facsimile bureaus. They are rapidly developing into permission-based, multimedia-based services bureaus (SMS, e-mail, voic-email, fax, etc.).

Notes

1. Currency figures are in Australian dollars unless otherwise noted.
2. See www.budde.com.au.

Multichannel Digital Broadcast Application over ADSL

Kenneth Clancy
Director, Technology Adoption
Aliant Telecom

Cindy Robichaud
Broadband Architect
Aliant Telecom

Kevin Hastings
Chief Architect, Digital Broadcast Services
Aliant Telecom

1. Abstract

This paper describes the development and deployment of a multichannel interactive digital television (iDTV) application.

In January 2000, Aliant Telecom launched a multichannel interactive digital television (TV) application over its broadband network. This application introduction resulted from a plan to migrate from providing "narrowband telephony services" to providing a full suite of interactive broadband multimedia services. This was deemed necessary to counter both the encroaching service capabilities of the cable and satellite providers and the potential customer losses. A full 6 megabits per second (Mbps) of bandwidth is required to the consumer residence over standard asymmetric digital subscriber line (ADSL) to meet the full interactive communications needs of the residence. This includes local and enhanced telephony, long distance, high-speed Internet (HSI) access, multichannel digital TV, and multichannel digital video to the personal computer (PC), all over a single copper wire to the home.

2. Telecom Industry Video Opportunity

2.1. The Competitive Imperative

iDTV is imperative to the success of incumbent telecommunications service providers (telecoms) facing intense competitive pressures, especially from cable TV companies. To date, two-way coaxial architectures as deployed by the cable TV providers and one-way satellite architectures as deployed by the direct-to-home providers have dominated the current broadband solutions to the home.

Previously leading in meeting consumer telecommunications needs, the telecom industry has generally been slow to embrace this new consumer demand. By waiting too long to deploy broadband networks and associated applications, telecoms missed their first opportunity, which allowed the cable companies to be first to market with HSI access service introductions. With cable companies now positioned to deliver a full suite of interactive broadband services such as local telephone, long distance, and HSI over their existing broadcast-based infrastructure, an immediate solution is required for the telecom that exploits the strength of its embedded copper investments and existing customer relationships.

Telecoms are now responding with fiber/copper solutions based on the deployment of ADSL. This paper will show how telecoms can leverage their broadband investments to deliver multichannel digital applications to the consumer.

2.2. Service Description

HSI and digital TV service are the telecom's anchor tenants on its high-bandwidth network. The following summary provides the current view of the iDTV service.

The service attributes include the following:

- 210 digital channels, comprising 123 broadcast video channels, 19 pay-per-view television channels, three Web channels, and 65 stereo audio channels
- Simultaneous support for two TV set-tops, TV–on–PC service, HSI access service, and local and enhanced telephony services
- Local on-demand video and audio content -
- An interactive program guide (IPG)
- Instant pay-per-view access on demand via the customer's remote control
- Internet access through the television; e.g., e-mail, Web browsing
- TV convergent Internet protocol (IP) portal with themed categories of content: news, weather, entertainment highlights
- Seamless navigation through the broadcast and Internet content
- 10 digital video broadcast channels to the PC
- Presentation of on-line or paper detailed bill for all

usage and subscription services: long distance, local, enhanced services, Internet, and iDTV

Pricing is competitive, and choosing a bundled service rewards customers with package discounts. For example, one service bundle offers the iDTV service plus unlimited use of each of the following: HSI, local and long-distance calling, and enhanced telephony. This bundling of services is a proven revenue protector, revenue generator, and, most importantly, positions the telecom for future revenue streams that will come from new broadband services and applications that use this same infrastructure investment (video on demand [VOD], e-commerce, and advertising).

3. Technical Solution

3.1. Technical Overview and Challenges

Figure 1 provides an overview of the integrated IP/asynchronous transfer mode (ATM) solution for delivering the multichannel digital broadcast application. In the digital head end, content is aggregated, groomed, and encoded into MPEG-2 or MPEG-4 IP multicast streams as appropriate for TV– or PC–based applications. Common telecom IP, ATM, and synchronous optical network (SONET) infrastructures are involved in broadcasting this content to broadband digital subscriber line access multiplexers (DSLAMs).

Distribution from the DSLAM involves a switched digital video architecture that streams channels to the customer's modem only when specifically selected by the customer's IPG (in the case of TV viewing) or PC client. Inside the home, a standard broadband ADSL modem connects to a PC or TV set-top equipment through standard Ethernet wiring solutions. The network must transport both multicast and unicast traffic as appropriate for broadcast and on-demand applications, respectively. Location of the head end can be centralized or distributed. This is an architectural choice that depends largely on the service provider's fiber transport, real estate, and operations assets.

The development and deployment of an iDTV application requires solutions to a number of technical challenges that are described in the following sections.

3.2. Video Encoding and Compression

The heart of the digital head end involves state-of-the-art satellite or fiber-based content acquisition systems and motion picture experts group (MPEG) compression systems. Pristine copies of all broadcast networks are required in order to generate high-quality compressed streams suitable for recovery at the digital set-top box. Provided that the sources are optimum, the process of encoding those programs to MPEG program streams will be the key determinant of picture quality.

The MPEG-2 encoders used for TV, and MPEG-4 encoders used in PC-TV applications accomplish fixed bit-rate stream generation for all video channels, their audio, and audio-only channels.

ADSL access networks, though broadband in nature, do not offer the "thick pipe" bandwidths seen in satellite or fiber to every customer; rather, the presentation of a 6 Mbps fixed pipe for all simultaneous services implies that bandwidth is fixed to reasonable allocations, despite the broadcaster's source bit-rate. This implies that in a two-TV environment, the MPEG compression process must be exceptionally efficient.

Sophisticated video noise-reduction technology is employed to minimize compression efforts for noise, and complex motion estimation processing is involved in maintaining high picture detail even in rapid motion content being broadcast.

FIGURE 1

Integrated IP/ATM Solution Overview

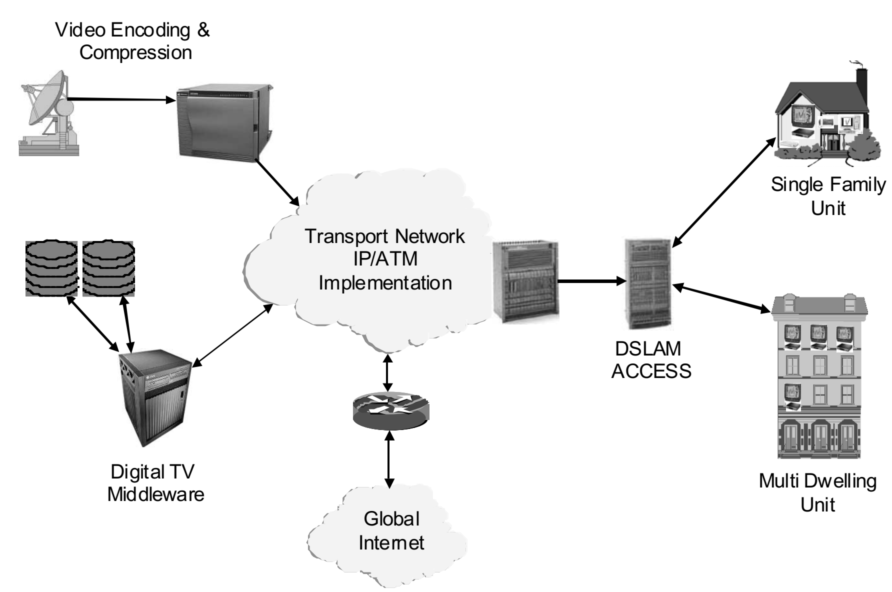

The competitive 7x24 nature of the broadcasting service dictates that the encoding facility be equipped with restoration facilities, spare equipment, and backup power systems.

Further, in-service maintenance maneuverability is assured through video-audio program switching equipment that interfaces the encoding systems and the satellite/fiber source systems. In regulatory environments where simultaneous program substitution is required to protect the distribution rights of local broadcasters, the video-audio program switch also meets this need.

3.3. Digital TV Middleware

A critical component of the solution is the digital TV middleware application. This application meets a number of critical requirements, including management of the user interface, customer subscriptions, and IPG and providing the operational service and interfaces required for customer activation, provisioning, and telecom billing.

The major capabilities provided by the middleware are as follows:

Content: The digital TV middleware manages the content availability in packaged format or on an individual channel subscription basis; i.e., pay-per-view (PPV).

Set-Top Management: The digital TV middleware manages the set-top and customer configuration and the software integration with the set-top's operating system. The set-top also translates the channel request to the network IP address required to source the appropriate content.

Interactive Program Guide: The IPG is a key navigation tool, given the large channel line-up involved. The middleware creates the IPG data fields and ensures that the information is delivered and available to the set-top according to consumer-acceptable channel-change requirements. Seven days of IPG formatted information is provided and managed.

Customer-Usage Reporting: The middleware also provides the event-management software, as well as the customer reporting. Application program interfaces (APIs) upload usage information so that the single bill for the packages and PPV events can be identified to the customer bill. User interfaces are provided to support all administration functions as well. Finally, broadcaster billing and invoicing is generated by this middleware to facilitate broadcaster content-distribution fee settlements.

3.4. Transport Network IP/ATM Implementation

There are several approaches that can be implemented to distribute digital content throughout a broadband network. One approach utilizes classical IP solutions and another is based on an integrated IP and ATM solution.

The classical IP approach, illustrated in *Figure 2*, utilizes Layer-3, IP–based technology. Constraints with the manageability and scalability of a pure IP–based multicast network become apparent as the number of customers and channels increase. Operations personnel would require intensive Layer-3 training in order to manage a complex IP multicast network where each individual channel must be carried as a separate entity. Also, signal integrity is not assured

FIGURE 2

Classical IP Distribution

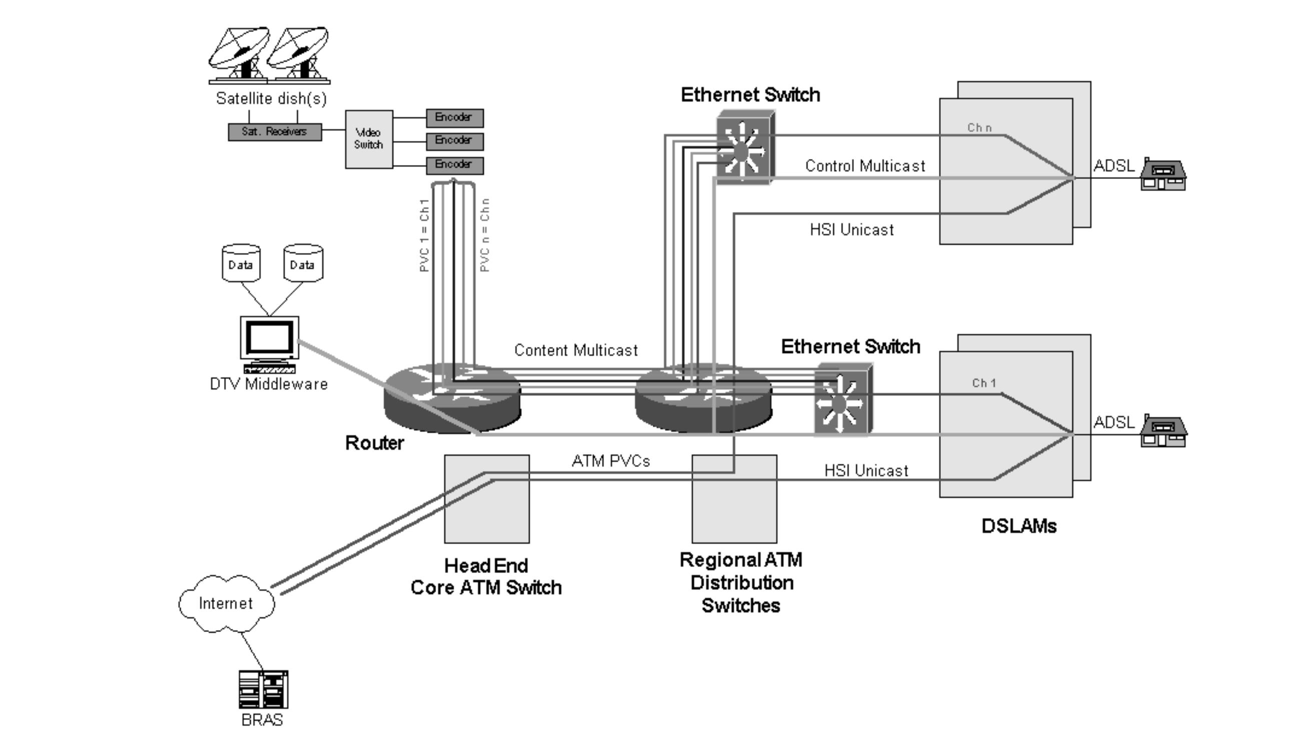

through current implementations of Layer-3 networks, as all traffic is of equal priority.

The second approach, based on an integrated IP and ATM design, offers several advantages in a telecom environment. ATM, by design, is a multiservice protocol. By implementing an ATM infrastructure, many of the complexities inherent in a Layer-3 infrastructure can be overcome. ATM enables the distribution of both unicast and multicast content on the same network using separate permanent virtual circuits (PVCs), as shown in *Figure 3*. Content can be inserted directly in to the ATM network, and since most DSLAMs are ATM–interfaced, distribution of the content to the customer is carried out with ease. The inherent quality of service (QoS) capabilities of an ATM network enables service providers to ensure that video traffic traverses the network without disruptive delays that could affect picture quality and channel-change latency.

The ability to format content as ATM traffic directly from a digital head end allows content to be collapsed onto point-to-multipoint ATM paths, which can easily be scaled and distributed throughout the entire network. The Layer-3 expertise required to manage complicated IP multicast traffic is no longer required by operations personnel who are already well trained in ATM technologies.

3.5. DSLAM Access Implementation

Several features and enhancements to improve broadband video distribution are becoming more prevalent in DSLAM vendors' offerings. A high-bandwidth link to the distribution network, such as optical carrier (OC)–12 ATM, allows the content to reside at the edge of the network where access times will be optimized. IP awareness within the DSLAM eliminates the need for a traditional Layer-3 overlay infrastructure and again optimizes content access speed. The network infrastructure necessary to support video applications and other feature-rich services simplifies, as the features available within the DSLAM improve. Higher-density shelves, with remote installation options, are also reducing the cost of creating a broader serviceable area.

Bandwidth requirements increase as services become more and more rich. Establishing a "service mask," or the tradeoff between bandwidth-rich services and servable areas becomes a key consideration, as illustrated in *Figure 4*. To maximize revenue opportunities, telecoms will have two service masks for their ADSL access network. In the example shown, the customer can be offered a full service broadband solution, including iDTV, His, and telephony at a 2 km loop length. Customers with loop lengths between 2 km and 4 km can receive both HSI and telephony services. Broadcast video on the PC is also available to any customer with HSI service.

Since interactive broadband services demand a high level of performance from ADSL, work must be carried out to achieve a stable connection to the customer's home. Testing to determine the optimum ATM and ADSL parameters to use on the DSLAM to ensure peak performance and stability of the ADSL line must be carried out. Careful consideration must be given, based on collected data, to the bandwidth available on certain copper loop lengths that contain

FIGURE 3

Multicast and Unicast Network

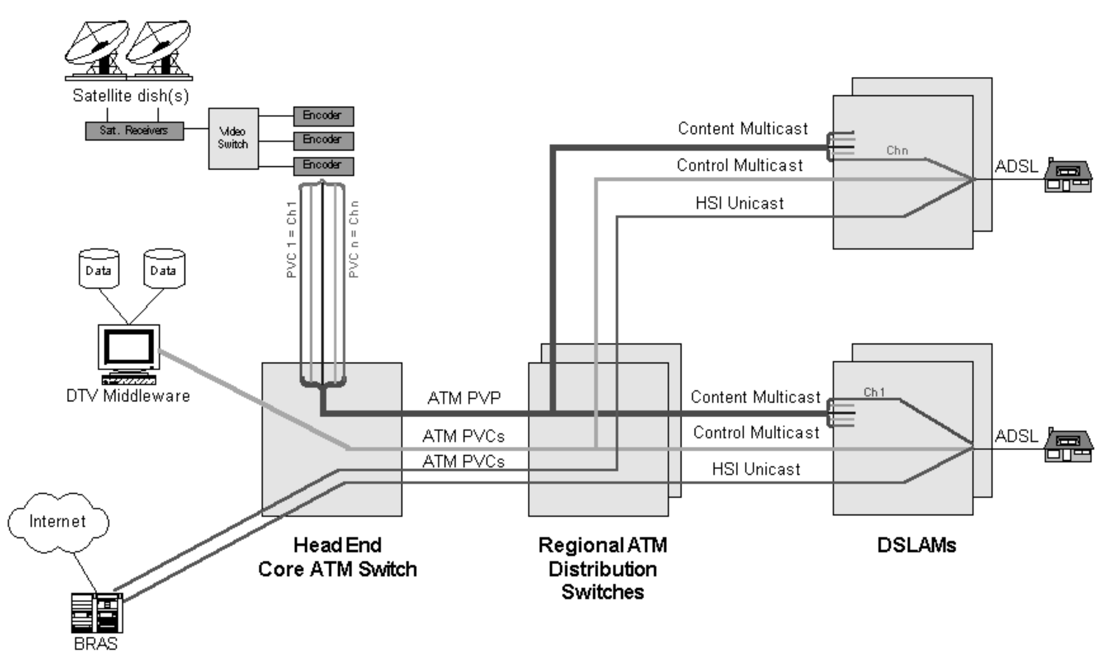

FIGURE 4

Service Mask

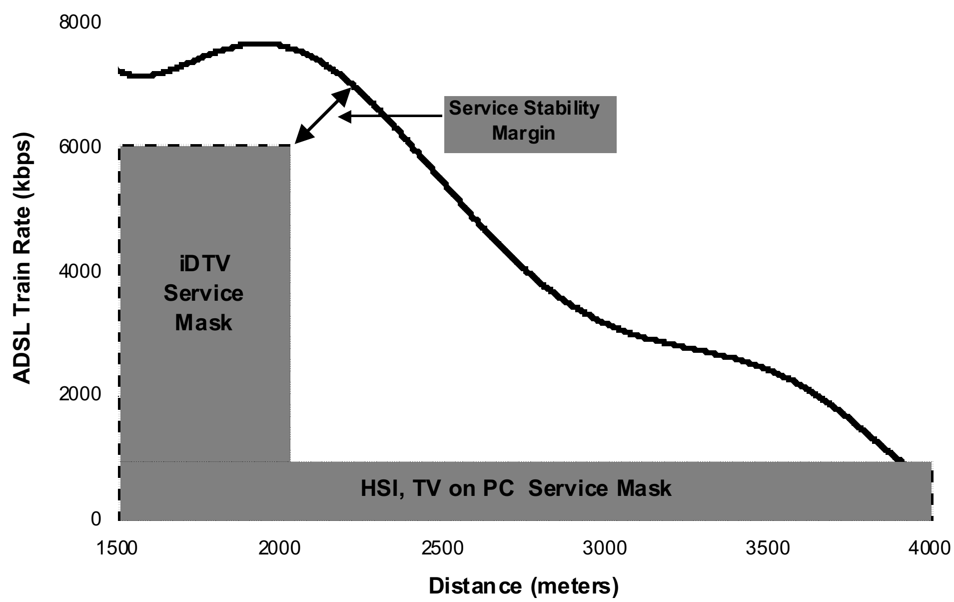

bridge taps and various gauges of copper wire, and guidelines for servable areas must be established. Technical improvements made by the DSLAM vendors have constantly been increasing the performance and stability of ADSL, which will improve the service readiness and reach of currently deployed infrastructures.

3.6. Residential Network

In the customer's home, new equipment must be introduced for iDTV service, including an ADSL splitter, ADSL modem, 10/100 BaseT hubs, TV set-top unit, and standard 10/100 BaseT wiring structure, as illustrated in *Figure 5*. The ADSL modem in the home is the key point of conversion from the carrier's ATM– and ADSL–based distribution network to the customer's 10/100 BaseT Ethernet–enabled home. The modem used must therefore be feature-rich and able to support multiple applications. This requires high throughput levels and advanced ATM PVC support that is becoming more readily available in standard HSI modems.

In approximately 50 percent of customer homes, the TV and PC appliances are located in the same room. Unfortunately, the remaining customers typically locate their two TVs and their PCs across several floors and rooms—making the modem and hub location a site-specific choice.

The popularity of 10/100 BaseT structures in modern homes, and frequent use of it in local-area network (LAN) and other PC environments, has made its introduction to consumers less difficult than initially expected. The growth of IP applications within the home is unprecedented and has contributed to application interworking challenges, such as dynamic host configuration protocol (DHCP) and security.

Finally, the cost of the customer-premises equipment (CPE) involved in this approach has been historically high and is only now descending toward an appropriate range, given the following:

- The emergence of digital set-tops incorporating the ADSL modem
- The widespread acceptance of 10/100 BaseT network interface cards (NICs) in store-bought PCs
- The development of customer-friendly self-serve applications for HSI access
- The mass marketing of consumer IP equipment (modems, hubs, routers), which has driven volume and fostered lower per-unit pricing

4. iDTV Results

4.1. Mandatory External Relationships

Relationships with at least two external agencies are required for successful iDTV service: appeasement of appropriate regulatory agencies and the establishment of appropriate business relationships with content owners.

Federal and municipal regulators have established relationships with incumbent broadcasters and may require time investments by new service providers to reach acceptance. This could range from federal broadcast license agreements to local permits or easements for Earth station and ADSL structures.

Telephone companies generally do not have relationships or credibility with most broadcast content owners, which introduces some risk to the broadband business where "content is everything." Federal broadcasting requirements, which must be respected, typically force the carriage of selected content, and even if that were not so, a competitive channel line-up normally dictates that formal content carriage contracts be signed with perhaps hundreds of content owners. In selected cases, these content owners may not

FIGURE 5

Residential Inside Wiring

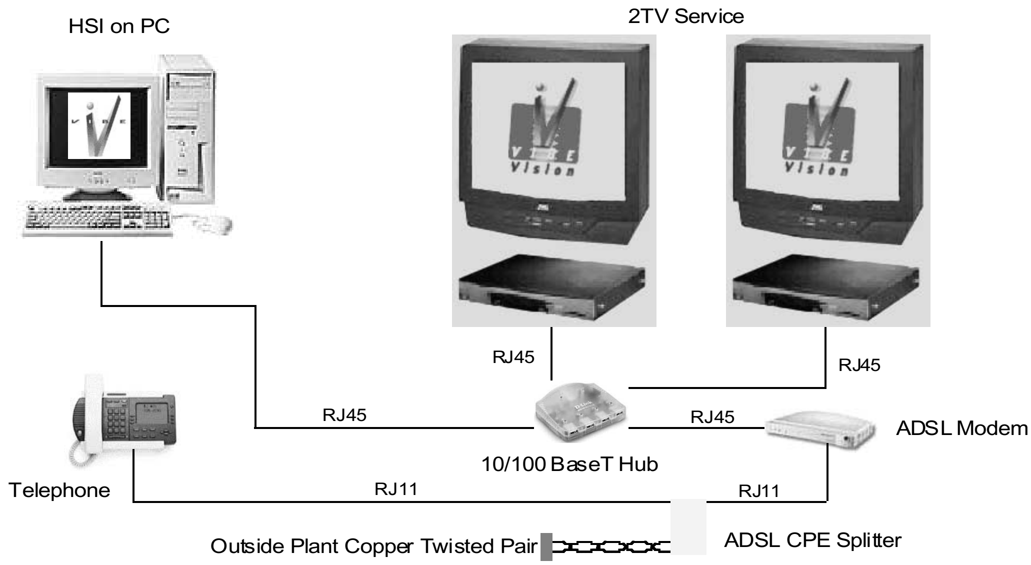

encourage new carriers, whose business presence and technology are unfamiliar to them. Program carriage fees paid to content owners represents the largest single cost to an established operator, and, as such, this topic has considerable impact on the iDTV business case.

4.2. Component Costs

There are new industry trends that have contributed to cost reductions during the past few years. The standardization of ADSL, the underlying technology developments in silicon, and the corresponding increase in volumes have reduced the DSLAM cost per port. Anticipating these trends, it then became urgent to focus on innovative exploitation of the supporting telecom technologies.

The major observations of industry trends are as follows:

- Component capital cost trends are progressing in the right direction: ~ 35 percent cost reduction on the total of these components over four years.
- Industry focus will drive further reductions, especially in the CPE cost area: set-top, modem, and wiring reuse.
- ATM/IP integration breakthrough occurred in 2000 and resulted in a $500 reduction per set-top served.
- High-density technology for the core and access components occurred in 2001, resulting in significant cost reduction and support cost reductions because of lower power and space requirements.
- Current design will support two TVs at 2000 meters.
- Cost does not include access rearrangement cost, which may be required in some jurisdictions.

The anticipated cost reductions have materialized as anticipated. *Figure 6* shows this trend over the past four years.

4.3. Revenue Opportunity

With technical trials now complete, the industry is now concentrating on enhancing the business model for the deployment of multichannel digital broadcast applications over ADSL. The initial deployment strategy is to focus on the telecom's most competitive major markets. The addition of iDTV to HSI will significantly increase the company's revenue per customer, as shown in *Figure 7*. When examining the early trends, the opportunities for increased revenues are clear.

Customers are currently adopting the new service in the home. The IPG is highly valued and has changed the customers approach to channel surfing. This is attributed to the interactive aspect of being able to get program information when desired. The use of PPV has also recently grown from 20 percent of the customer base to 52 percent. Again, this growth is attributed to the customer's ability to instantly select a movie or pre-schedule a movie, all from the convenience of their home.

5. Conclusion

Telecoms must respect the consumer's need for converged services. Now is the time to deliver the applications that will meet the consumer expectations and generate new revenue for the telecom. Those service providers that do not embrace this trend will face higher opportunity costs, continued market-share erosion, lost competitive advantage, and revenue loss.

This paper has shown the current viable technical solutions available to the telecom for delivering a full service application set to the consumer, including multichannel digital

FIGURE 6

Component Costs Trends

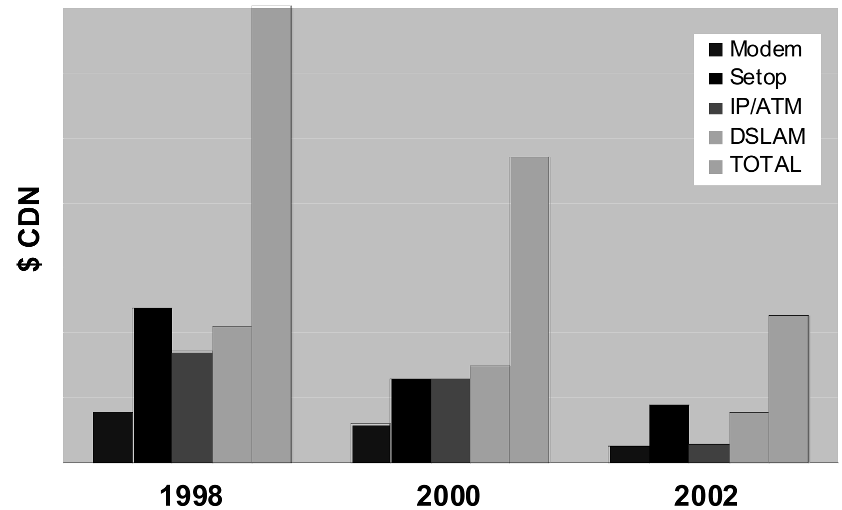

broadcast. The challenge remaining to the telecom is to recognize the opportunity, adopt an architecture, and complete a business model.

6. Further Reading

- Homepage for Aliant Inc.; www.aliant.ca
- Homepage for Aliant's VibeVision™ service; www.vibevision.com
- "Challenges and Opportunities of Delivering IP–Based Residential Television Service," *IEEE communications Magazine*, Vol. 38, No. 11 (Hastings and Nechita, November 2000), http://www.wtc2000.org/pdf/ab_088s.pdf
- Cindy Robichaud, "ADSL Access Requirements for Interactive Multi-Media Networks," DSLcon Toronto Conference, May 2000.

FIGURE 7

Monthly Broadband Revenue Opportunity

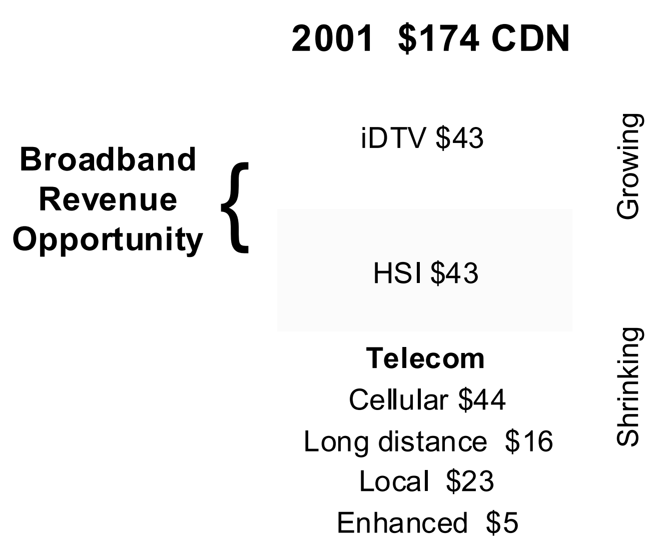

The Benefits of Automated Decision Analytics

Jeffrey Cleveland
Senior Principal and Director, Credit and Collections Risk Management Practice, New Media and Communications
American Management Systems, Inc.

Overview

This position paper details the following:

Two *primary challenges* faced by telecommunications providers today in the management of credit risk:

1. The assessment of up-front credit risk during the acquisition and provisioning process
2. The ongoing assessment of risk in the accounts receivable portfolio over time

A potential *solution framework*, which includes the following components:

- Up-front credit scoring and automation
- Periodic rescoring of accounts using behavior scores
- Automated decisioning capabilities
- A data mart for capturing, retaining, and analyzing account performance

The anticipated benefits stream, obtained from the following:

- Increased revenue
- Reduced Net Bad Debt (NBD) expense
- Improvement in cash flow (as measured by DSO)
- Reduced operating expenses

Current Challenges

In the telecommunications industry today, few firms have adequate processes and policies for measuring and managing credit risk at the customer or account level. And in those institutions that do happen to employ various risk-management strategies and tactics, there is more often than not the fact that the desired levels of benefits have not been met. In general, these shortcomings can be viewed as a result of two factors:

1. A lack of understanding and management around initial credit risk for new applicants
2. The inability to measure the ongoing riskiness, and changes in those levels of risk, in the overall customer portfolio

Up-Front Credit Assessment

Inconsistency in Decisioning: Decisioning data, provided by the applicant and obtained from third parties (i.e., credit-reporting agencies), is used to determine likely credit risk, overall creditworthiness, and applicable risk mitigation strategies and tactics. However, since these credit evaluations are often performed on a subjective and manual basis, most telecommunication firms' credit-provisioning standards are subject to a variety of individual standards, criteria, and biases. These factors result in a decisioning environment that is, from a credit-granting perspective, inconsistent. This inconsistency, in turn, subsequently leads to the following:

- Difficulties in measuring credit-decision effectiveness
- Potential increases in overall portfolio risk due to provisioning of higher-risk accounts
- Potential suppression of revenue due to declination or application of risk mitigation strategies and tactics to accounts that, in fact, do not present any undue incremental risk to the firm's portfolio

Inefficiency in Initial/Pre-Screen Processing: At many telecommunication firms serving the small- and medium-sized–business markets, there is often a major disconnect between the sales, marketing, and risk-management staffs. For example, during the initial sales evaluation phases, the sales staff frequently lacks any kind of tool to quickly pre-screen or pre-qualify, from a credit-risk perspective, potential or target accounts. This results in process inefficiencies at the sales level, as business developers may be spending time trying to target and acquire potential customers whose risk profile places them outside of the firm's acceptable and/or manageable levels of risk.

Ongoing Portfolio Assessment

Once a telecommunications firm has provisioned an account on their network/system, there is typically no set method or process to continually re-evaluate the overall risk composition of the accounts receivable portfolio. Specific challenges include the following:

Inability to Measure and Understand Changes in Portfolio Composition and Risk: Customers' levels of risk over time tend to be dynamic rather than static. That is, accounts that were viewed as "low risk" when they were part of the through-the-door population at time of acquisition and provisioning may very well turn out to be higher-risk customers. Likewise, through-the-door customers that were originally deemed "high- risk" might, by their own performance and behavior, actually turn out to be fairly low-risk accounts. The key element required in proactive portfolio management is the ability to measure and understand the following:

- The *distribution* of portfolio credit risk at any given point in time
- The *changes in distribution* of portfolio credit risk over time

Inability to Respond to Changes in Portfolio Risk: It is highly likely that a customer's level of credit risk that is assigned during the up-front credit assessment process is the same level that remains assigned to the account during that account's life. This poses two serious risks:

1. Accounts whose behavior indicates a lower level of risk may still have revenue-suppressing risk mitigation strategies and tactics being applied to it.
2. Accounts whose behavior indicates a higher level of risk may be avoiding the application of risk mitigation strategies and tactics being applied to it.

Inability to Measure Cost-Efficiency and Cost-Effectiveness of Strategic and Tactical Decisions: Over the lifetime of a given customer or account, hundreds of decisions are made, with each decision intending to affect customer behavior in some way, shape, or form. However, most firms have no way to determine which individual decisions are having the most impact and influence on customer behavior without a systemic way of continually testing and evaluating decision effectiveness.

Potential Solution Framework

Addressing these challenges in managing credit risk, both during the initial credit assessment and over the course of an account's life, requires a multiphased approach:

- Up-front credit scoring and automation
- Periodic re-scoring of accounts using behavior scores
- Automated decisioning capabilities
- A data mart for capturing, retaining, and analyzing account performance

Up-Front Credit Scoring and Automation

To address the issues regarding inconsistency in decisioning and inefficiency in pre-screening during the initial credit process, telecommunications firms can take their cue from the financial services industry, which makes extensive use of statistically derived *credit models* to predict initial/early life credit risk.

Credit models use applicant performance data as reported by *other* credit grantors and are used to assess the likelihood of risk in the absence of any direct experience with an applicant. Basically, the use of a credit score allows potential credit grantors to answer the following question:

"Since my firm doesn't have any direct performance experience with this applicant, how likely is this applicant to become seriously delinquent and/or write-off with us based on this applicant's credit usage and payment performance with other creditors?"

The model life cycle has four distinct phases:

1. *Model Strategy Design:* It is crucial in the beginning of any model development/deployment effort to outline the high-level designs of the overall program. For example, what type of behavior and results are attempting to be predicted? Where in the overall process will the model(s) be introduced and calculated? What kind of infrastructure is required to implement and support the model(s)?
2. *Model Development:* This phase typically involves a great deal of data preparation, from data access to cleansing and manipulation. Once the data is ready, the actual processes involved in creating and validating the scorecard are relatively straightforward.
3. *Model Implementation:* Once the model has been developed, the model scorecard must be integrated into and tested with existing production environments.
4. *Model Usage:* It is typically at this point where statistically trained model developers exit and more broad experience-based portfolio managers come into play. Understanding how to use models, particularly in areas such as time series analysis, score cut-offs, and controlled experimentation, is critical at this stage to gain the maximum leverage and lift from the overall modeling effort.

In addition to the use of credit models, the Internet makes it possible to deploy a Web-based credit system for use by both credit analysts and business developers. By utilizing an automated tool for data capture, scoring, and decisioning, telecommunications firms can significantly improve the consistency in credit decision-making, as well as improve the overall throughput in the credit process.

Periodic Re-Scoring of Accounts Using Behavior Scores

The use of credit scores on the through-the-door population is designed to predict short-term credit risk (i.e., within one year of account acquisition and usually inside of six months). However, once an account has been active for four to six months, behavior and performance are much better overall predictors of future risk. More importantly, that customer's level of risk changes as their behavior changes over time.

To measure these changes in riskiness, successful companies deploy statistically derived *behavior models* to predict ongoing credit risk for accounts in their portfolio.

Behavior models use a firm's own customer performance data to assess the likelihood of future risk based on past performance. Basically, the use of a behavior score would allow a telecommunications firm to answer the following question:

"Based on my firm's direct experience with this customer, how likely is this customer to become seriously delinquent and/or write-off some time in the future, based on this customer's credit usage and payment performance?"

The same four model life-cycle phases described in the previous section are also applicable here: Model Strategy Design, Model Development, Model Implementation, and Model Usage.

Automated Decisioning

When using predictive models, such as credit scores or behavior scores, to assess potential future risk, one of the two significant considerations to be taken into account is at which point to set the score cut-off. For example, when using a score that has a range of 0 to 500 (where the lower the number, the higher the risk, and vice versa), a decision may be made to set a score cut-off at, say, 300. That decision drives two actions: accounts scoring above 300 may have no action taken against them, while accounts scoring at or below 300 would have a certain treatment applied.

The second important consideration is that of the treatment itself, which is, in theory, designed to influence account behavior.

Like all decisions, however, these two types are subject to specific scrutiny:

"How do I know if the score cut-off set by my firm is at the right point? Perhaps it should be higher? Or lower?"

"Once the scoring decision has been made, how do I know if the treatment action is the most cost-efficient and cost-effective one available?"

The way to determine the answers to these types of questions comes from continuous learning through ongoing decision testing. For example, to test the cost-efficiency and cost-effectiveness of treatments, accounts that are statistically similar would be segmented randomly, with some accounts receiving "Treatment "X" and other accounts receiving "Treatment "Y." Both decision and subsequent account performance data (including delinquency/NBD measures and treatment costs) would be captured over time and analyzed. This type of testing must be performed in a production environment.

Leading firms (again, looking at the financial services industry as a guide) use an automated decision engine to implement and manage this type of testing in their production environment. An automated decision engine would need to contain the following characteristics to provide the right benefits in terms of managing credit risk:

- Customer assessment using predictive model scorecards (both credit models and behavior models)
- Customer segmentation of accounts into logical groupings
- Flexible decision logic aimed at providing differentiated treatment
- "Test-and-learn" capabilities to continually challenge assumptions about how decisions are (or aren't) affecting customer behavior
- Volume-simulation processing using production volumes to determine operational impacts of potential decision changes
- Comprehensive reporting, including decision effectiveness reporting, model validation, and score/attribute reporting

Account-Performance Data Mart

One of the key elements in any credit risk-management program is the analysis of data. More than just "standard reporting," detailed time series analyses of account performance and decision data is the hallmark of leading-edge accounts receivable management firms, particularly within the financial services industry.

There are four types of data that need to be captured and retained for analysis purposes:

1. *New Account Data:* This includes all relevant data on the through-the-door population in general and on provisioned new accounts in particular. Basically, how did the applicant look when they first came through the door, and how were they evaluated and treated?
2. *Monthly Performance Data:* This type of data, usually obtained from the billing system(s) on a monthly basis, is where detailed account performance, such as usage, payments, and newly calculated behavior scores, is obtained.
3. *Decision Data:* At each and every decision point that affects customer behavior, a copy of that decision, along with the account in question, is stored for future analysis.
4. *Treatment Data:* This includes any type of risk mitigation strategies and tactics, and collection treatments.

In addition to the capture and retention of data, a method of accessing and manipulating the data is required as well.

Benefits

By implementing the previous four initiatives, telecommunications firms can expect to see both top-line and bottom-line benefits in the following areas:

- Increased revenue
- Reduced net bad debt (NBD) expense
- Improvement in cash flow (as measured by days sales outstanding [DSO])
- Reduced operating expenses

Increased Revenues

By using predictive scores (both credit and behavior), companies can better segment their customer base throughout the entire credit life cycle, from account acquisition and provisioning to balance management, to collections and recoveries.

Since predictive scores allow for the rank ordering of accounts by risk, telecommunications firms can avoid treating low-risk accounts (i.e., those accounts with a low probability of becoming delinquent and/or writing off in the future) with tactics and actions aimed at higher-risk customers. By treating low-risk accounts with, say, non-aggressive treatments, it is possible to avoid the potential (and unintentional) suppression of revenues that might not have been permitted by low-risk accounts being treated with more aggressive tactics. Additionally, there may be improvements in customer retention.

Moreover, by deploying automated decisioning capabilities and a data mart to continually test and measure the efficiency and effectiveness of strategy and tactics decisions, many firms are able to quantitatively prove that they are optimizing their decision-management framework.

Reduced NBD Expense

In most instances, *reductions in NBD represent the greatest opportunity for bottom-line improvement*. Predictive scores (both credit and behavior) provide companies with the ability to divide both their applicant and customer bases into risk-based segments and match treatments (i.e., aggressive versus lenient) with associated risk classes.

Reductions in NBD come from both sides of the write-off equation. First, telecommunication firms experience a *reduction in the number of accounts that become delinquent and go to write-off*. This is achieved primarily through the use of credit scores to exclude high-risk accounts (that is, accounts that have a higher propensity to become delinquent and/or write-off than the firm is willing to accept) during the new account acquisition and provisioning process.

Second, companies experience a *reduction in the average write-off balance for accounts that do end up writing off*. This is achieved primarily through the use of behavior scores to minimize balance build-up on accounts that, although non-delinquent today, are at serious risk of becoming delinquent and/or writing off in the future.

Additionally, by deploying automated decisioning capabilities and a data mart to continually test and measure the efficiency and effectiveness of strategy and tactics decisions, it is possible to quantitatively prove that the decision-management framework is being optimizing.

Improved Cash Flow (DSO)

By improving their ability to manage both the number of delinquent accounts and their balances, telecom providers' integrated risk-management approach acts as a force multiplier: In addition to reductions in bad debt, improvements in cash flow as measured by DSO are often achieved.

Since DSO is essentially a measure of the distribution (by age) of the accounts receivable portfolio, a successful risk-management program results in fewer delinquent accounts and lower average balances on those accounts. The net result is that fewer A/R dollars (because of fewer accounts and lower balances) end up being delinquent (i.e., skewed toward the "right side" of a distribution curve). Conversely, more A/R dollars, in absolute and percentage terms, end up being current (i.e., skewed toward the "left side" of a distribution curve). This translates into improved cash flow.

Reduced Operating Expenses

Most companies experience savings in expenses associated with the management of their credit-risk processes, primarily through reduced account volumes in collections. This is achieved in two ways. First, the use of up-front credit scores during the new account acquisition and provisioning process allows firms to exclude high-risk accounts that would most likely end up in collections.

Second, the use of behavior scores, particularly those developed for newly delinquent accounts, provides for the rank ordering of accounts in collections by risk. These rank-ordered accounts are then matched against risk-differentiated treatment strategies, allowing for limited collection resources to be applied to accounts in the most need of direct collector contact. Conversely, lower-risk accounts (i.e., those accounts most likely to pay with minimal or no collector intervention) are then routed to collection-treatment scenarios that don't incur higher-costing collector intervention.

Other Benefits

Many telecom providers are able to achieve two other benefits that, while not directly bottom-line–measurable, are crucial to effective credit-risk management:

Consistency in Decisioning: In most manual/subjective environments, the application of decisions is inconsistent across two compounding dimensions: 1) different individuals make different decisions, and 2) the same individual makes different decisions over time. Leading-edge companies resolve these two inconsistencies through the use of predictive scores (both credit and behavior) and the application of decision logic through an automated decision engine.

Ability to Prove Decision Optimization: By incorporating "test-and-learn" strategies across the entire decision-management framework in conjunction with a dedicated data mart, both "champion/challenger" and "control" groups can be created to prove quantitatively that all credit-risk–related decisions are being made to optimize overall A/R portfolio results.

FCC Tackles Number Conservation—Again

David Crowe
President
Cellular Networking Perspectives Ltd.

In December 2000, the Federal Communications Commission (FCC) released a second Report & Order[1] regarding numbering resource optimization—the conservation of the rapidly dwindling North American numbering resource. Although the FCC only has authority over U.S. phone numbers, and even though it shares this authority with the States, any changes adopted by the FCC will have a big influence on Canada and the Caribbean nations that share this numbering space.

Wireless carriers will be heavily affected by the various initiatives being considered by the FCC, particularly number pooling.

Thousands-Block Pooling

The FCC is still enamored of thousands-block pooling, implying it is superior to other methods of number conservation. Unfortunately, number pooling works only if wireless carriers adhere to the same number of assignment boundaries (rate centers) as landline carriers, and this will reduce the efficiency of number assignment considerably.

Thousands-block pooling allows the current basic number allocation unit (10,000 numbers) to be shared among as many as 10 carriers. This could also be done by moving from a 6-digit to a 7-digit number analysis, but this has not been considered by the FCC, although it may have less impact on switching equipment.

It might seem unnecessary to share blocks of 10,000 numbers among carriers when there are hundreds of millions of phone numbers in use. There are two reasons for this:

1. Competition in the local-exchange market is resulting in more companies assigning numbers to phones. Note that competition in the long-distance market did not have this effect because these carriers do not assign phone numbers.
2. The rate centers are small. The rate center is the geographical area within which numbers are assigned by local-exchange carriers (LECs). This means even large customers may not always be able to fill a block of 10,000 numbers.

Competition is only going to get more intense, but rate centers are artificial areas that could be merged together, allowing more carriers to obtain several thousand customers in each rate center, which would make pooling unnecessary.

The reason behind the FCC's enthusiasm may be that pooling makes use of the local number portability (LNP) infrastructure mandated by them, and it therefore may make the expense and complexity more palatable. Number pooling, however, has some important differences from LNP. Whereas LNP has to support the migration of single or small blocks of numbers to another carrier, number pooling starts with at least 1,000 numbers moving. Number-portability systems (NPAC Release 3.0) have had to be modified to represent a pooled block by one record instead of 1,000. Straightforward implementations of LNP would collapse under the data storage requirements of number pooling.

Wireless Carriers and Pooling

Wireless carriers cannot participate in number pooling unless they adapt their number-assignment procedures to those of rate centers. Wireless carriers generally assign numbers in larger areas, resulting in a larger local calling area. If wireless carriers were to conform to rate center boundaries, their number allocations would instantly become much less efficient, and number pooling cannot be guaranteed to return wireless carriers to the level at which they started!

Timing for Wireless Carriers

The FCC has ruled that wireless carriers must support number pooling as soon as they are able to support LNP. The FCC declined to give any transition interval, even though LNP and number pooling are not identical capabilities.

Rate-Center Consolidation

Rate centers are the fundamental unit of number assignment and billing in the North American Numbering Plan (NANP). Every number within a rate center is considered to be in the same location for billing purposes, meaning that tolls cannot be charged within a rate center.

In many cases, there are multiple, adjacent rate centers between which long-distance charges do not apply. These could be consolidated with minimal impact on billing. By doing this, all number blocks within the enlarged rate center would be available for assignment. The only potentially negative effect of this is that the calculation of charges based on distance would be slightly different—higher for some calls and lower for others. However, with distance-insensitive billing for long distance increasing in popularity, and with the differences being small, this is probably not a serious problem.

Combining rate centers between which toll charges apply is more controversial; however, any loss here could easily be handled by adjusting local rates such that consumers would not see an increase in the overall costs of telecom services.

The FCC is warming toward rate-center consolidation. It would probably be more enthusiastic if rate-center boundaries were under FCC control, rather than under state control. Number pooling, by contrast, is largely within FCC jurisdiction.

Thresholds

The FCC requires carriers adhere to number utilization thresholds before being eligible for more numbering resources. The threshold will be set at 60% on March 29, 2001, and this will increase by 5% on June 30, 2002, and by another 5% each following year, to a maximum of 75% on June 30, 2004.

The calculation of thresholds is strictly:

$$\frac{\text{Numbers assigned to customers}}{\text{Numbers assigned to carrier}}$$

Numbers used for testing, numbers used for administrative purposes, or numbers being aged (i.e., taken from one customer but not eligible for reassignment to another customer for some time to prevent confusion) are not counted as "assigned to a customer." Thresholds are calculated separately for each rate center.

Audits

The FCC has claimed the right to audit carriers, randomly or for cause, to determine what their use of numbers really is.

Assigning New Area Codes: Splits versus Overlays

When numbering resources within an area code are nearly exhausted, a new area code must be assigned. To maintain number conservation objectives, this should be done efficiently. The basic choices are as follows:

- Splitting: Assigning a new area code to some existing customers
- Overlaying: Assigning a new area code for new numbers only

Area-Code Splits

Splitting an area code is the traditional way to assign new numbering resources, but this results in half of existing phones being assigned new numbers. Seven-digit dialing between the area codes can be preserved by "protecting" certain office codes from assignment in the new area code, but this does not meet number conservation objectives. Ten-digit dialing from one area code to another will be necessary, eventually, to allow all office codes to be utilized.

Figure 1 illustrates how splits can be performed. Often, the suburban area will receive the new area code (e.g., as in Dallas and Atlanta), but it is possible to split a region down the middle (e.g., a city straddling a river) or in any other way dividing the population approximately in half, with some consideration given to expected growth patterns.

FIGURE 1

Area-Code Splits and Overlays

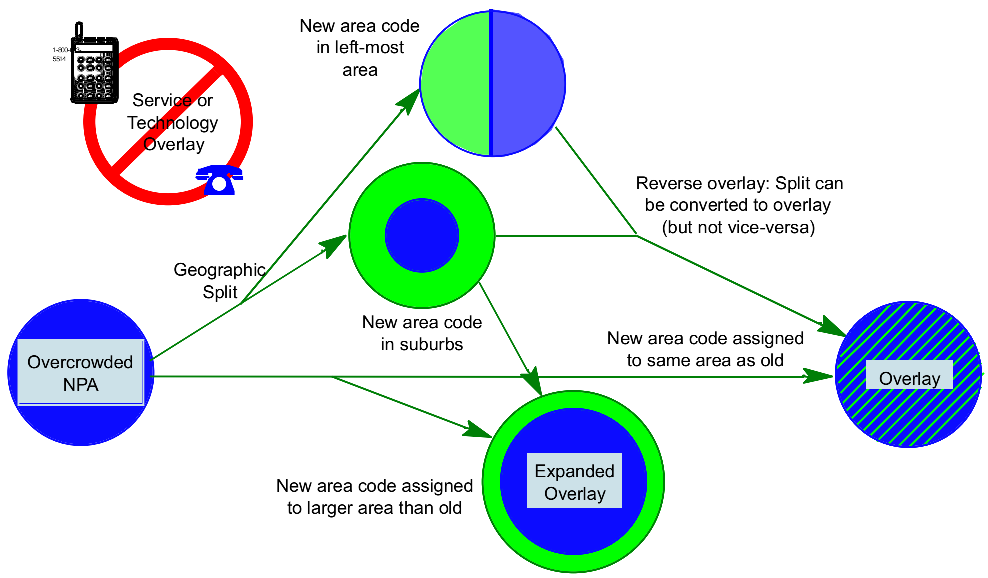

Overlays
Overlays assign the new area code to the same geographical area as the old area code. The FCC has banned the use of service-specific or technology-specific overlays (e.g., all wireless phones in the new area code) as discriminatory. Overlays are used, instead, for new numbering requirements. This puts newer customers at a disadvantage, because the majority of customers need to dial 10 digits to reach them (and vice versa). Consequently, mandatory 10-digit dialing may be applied so that everyone may suffer for the benefit of the minority.

A variant on the overlay is the *expanded overlay*. In this instance, the new area code applies to a larger geographical area. This allows one area code's resources to provide relief to two existing areas, which is useful when growth is expected in suburban areas. Another variant, a *reverse overlay*, allows the conversion of an area-code split into an overlay; the assignment of both area codes is allowed over the entire area—a useful approach when population growth projections proved to be invalid.

No Nationwide 10-Digit Dialing

The FCC will not force nationwide 10-digit dialing, but it will creep in as area codes get smaller with splits and as areas affected by overlays mandate 10-digit dialing.

"D" Digit Expansion

The fourth digit of a NANP number (when expressed in its 10-digit format), known as the "D" digit, is restricted to "2" through "9" to allow 7-digit dialing. If this digit was a "0" or a "1," some switches would be confused, not being able to differentiate between local calling and operator/long-distance services.

Expanding the "D" digit would expand the capacity of each area code by about 25 percent, but nationwide 10-digit dialing would be a prerequisite. The FCC decided not to mandate this step—yet.

Conclusions

The FCC still favors thousands-block pooling, even though this may waste as many numbers as it saves, will cause wireless carriers significant costs, and could cause them to fail to assign numbers as efficiently as they do today.

Table 1 compares the major forms of number expansion and conservation being promoted by the FCC. Splits and overlays are required when numbering resources are exhausted and when the choice between them is not straightforward. Number pooling is inferior to rate centers because of its reliance on LNP, with its cost and reliability issues, and because carriers will be forced to adhere to rate-center boundaries, which, unless rate-center consolidation occurs, cannot be done efficiently. Additionally, if rate-center consolidation does occur, the need to share pools of less than 10,000 numbers may not exist.

Notes

1. The Report & Order is available at www.fcc.gov/Daily_Releases/Daily_Business/2000/db1229/ fcc00429.doc (or .txt).

TABLE 1

Comparison of Number Conservation and Expansion Methods

	Does not require LNP?	Wireless carriers can define local calling areas?	Customers keep current numbers?	Optimized use of existing resource?	Preserves 7-digit dialing?
Area-Code Split	√	√			Partially
Overlay	√	√	√		
Number Pooling			√	√	√
Rate-Center Consolidation	√	√	√	√	√

Revitalizing Legacy Business Support Systems

Rolando Espinosa
Senior Vice President, Products and Strategy
Abiliti Solutions

It has been five years since the Telecommunications Act of 1996 legislation created this marketplace. A lot of "noise" has been produced by the changes in the investment in and valuation of technology companies, but through this, a very traditional trend has developed—the maturing of an industry.

As with other telecom markets that preceded it, the measure of a competitive local-exchange carrier's (CLEC's) value has shifted over time. Initially, network deployment and geographic coverage drove valuation. In other words, does the CLEC have the potential to succeed?

Second came customer acquisition. Could the company demonstrate its ability to succeed by winning customers?

Today, the CLEC market is in its third phase of maturity, where success is no longer measured by size alone, but by profitability.

A CLEC's view of its investment in business support systems (BSSs) needs to change to keep pace with its evolving objectives. In the early life-cycle phase, a BSS application was a security blanket facilitating growth to assure investors and customers of its stability. The more a CLEC spent on BSS, the greater the security.

In today's environment, efficiency is the goal of BSS expenditures. A dollar spent in BSS infrastructure must facilitate a cost savings or revenue increase of more than a dollar.

A shift from organic growth to acquisition creates challenges for the BSS infrastructure.

Firms that cannot achieve profitability become takeover candidates. Strong CLECs are becoming both larger and more profitable by acquiring their weaker brethren. A potentially negative by-product of this otherwise sound course of action is that CLECs are left with multiple BSS applications to manage and, therefore, face increased expenses.

- ***Training Personnel on Multiple Systems***
  Acquisitions create the redistribution of employees and ultimately layoffs. Employees need training to work on systems with which they are not familiar and have distinctly different user interfaces. Worse yet, employees may be required to work on multiple systems simultaneously. The efficiencies and profitability envisioned by the acquisition are then difficult to achieve.
- ***Implementing Products and Services in Multiple Systems***
  Acquisitions often result in the overlapping of markets and products and redundant BSS applications, which leads to the need for duplicate data entry. If a carrier wants to offer a common product suite across markets, this will usually require dual entry into multiple BSS applications. This also assumes that the two systems can support the same product set and service definitions.
- ***Developing Interfaces to New Call Record and Event Formats in Each System***
  Acquisitions result in new networks or network components or new downstream BSS components using the same network. Either way, if the network changes, through the acquisition of new elements or the introduction of new features or services, this will affect the interface/handoff of call detail records (CDRs) to downstream BSS components. A greater number of standalone components leads to a greater number of changes required and, therefore, increased cost.
- ***Managing the Same Information in Disparate Systems***
  Having the same customer in two systems (for example, one for Internet protocol (IP) billing and one for voice) requires dual entry. This translates to twice the time and potentially increases the number of mistakes. The more systems a CLEC manages, the greater the time and effort to maintain consistency. Maintenance costs increase exponentially with each new system addition.
- ***Producing Multiple Reports from Each System***
  Not only are users forced to interact with widely different user interfaces, but the reporting methodologies and tools may be vastly different as well. Information is fragmented across multiple systems, and the cost to converge databases and reports is another drain on the profitability of a CLEC.

The end result is inefficiency and revenue leakage due to redundancies, omissions, and inconsistencies.

The rollout of new products and services present another challenge to the BSS infrastructure.

When CLECs first entered the market, customers were looking for any alternative to the incumbent local-exchange carriers (ILECs). This demand coupled with the ability to bundle local and long-distance services and under-cut ILEC pricing plans fueled customer acquisition. However, this flood of customers will not last forever. New products and services are necessary to continue growth.

New service offerings lead to increased profitability for two reasons. First, the cost of upselling a new product to an existing customer is significantly less expensive than acquiring new users. Second, customers using a wide variety of a CLEC's services are less likely to churn.

CLECs are currently introducing the following products and services:

- ***Usage-Based IP Services:*** This includes billing customers based on actual usage, such as MB of disk space, e-mail messages sent, session duration, bytes transferred, number of documents downloaded, online gaming, etc.
- ***Quality of Service (QoS)–Based Pricing:*** This builds QoS metrics (e.g., network availability, connection setup time, and error performance) into pricing plans or service-level agreement (SLA) credits.
- ***Content Services:*** This includes the delivery of secure data transmissions, streaming media, frequent updates, or mission-critical applications.

How best to address the BSS challenges?

Whether facing a legacy BSS unable to support new products or determining how to manage a BSS environment, carriers have three options:

- ***Replace Their BSS Applications or Consolidate to One System.***
  This is the commonly implemented solution. However, the complexity of the migration increases with carrier size. Integrated links established over time with other operations support system (OSS) applications must be simultaneously redeveloped. Many employees must be retrained. During a large migration, a carrier will be less responsive to customer needs and competitive actions, posing a great detriment to a CLEC's success over the long term.
- ***Support New Services by Implementing Additional BSS Applications, Consequently Creating a Multiple Systems Environment.***
  There are inherent inefficiencies maintaining multiple systems. The expense of a secondary BSS (whether it supports a narrow product set or small geographic territory) will be allocated against a smaller revenue stream.
- ***Add an Event-Management and Rating System to the Legacy Environment to Overcome Limitations.***
  This solution enables a carrier to gain the necessary flexibility to address the unique needs of new services. This approach also offers more efficient transmission of data both into and between BSS applications, while avoiding costly duplication.

As the market matures, investment-driven growth must be replaced by profit-driven growth. CLECs today must be able to survive solely on the revenue they can garner from product and service sales. For this reason, greater emphasis must be placed on "revenue mining." Service providers must ensure they are mining their networks for every last billable nugget they can find. They must ask the following:

- Am I getting revenue out of each transaction on my network?
- How do I know I'm getting the maximum revenue available?
- Am I getting the most out of my existing legacy systems, or am I spending hard-earned dollars repeating the same tasks on multiple systems?
- How can I ensure that I'm billing all I can and that I am doing it at the lowest possible cost?

CLEC's can use a single, convergent event-management and rating system in conjunction with their legacy BSS applications to increase revenue from new service introductions while maintaining efficient, cost-effective operations.

This paper was previously published at CLEC.com, September 1, 2001.

Changing Business and Its Paradigms: Post 9/11/01

Robert S. Fleming
Professor, Management and Management Information Systems
Rowan University

The world in which we live, work, and travel changed on September 11, 2001. The terrorist attacks of that day and the ensuing attacks involving Anthrax letters following that date have caused individuals, groups, and organizations around the world to ponder the new realities of terrorism and attempt to make sense out of these new developments.

Unfortunately, the reality is that vulnerability to acts of terrorism has been recognized for some time. The 1993 bombing of the World Trade Center, the 1995 bombing of the Alfred P. Murrah Federal Building in Oklahoma City, and a number of less visible and publicized smaller events have called attention to this vulnerability well in advance of September 11, 2001. Those responsible for the information technology (IT) function in contemporary organizations have recognized this threat, particularly as it relates to cyberterrorism, for some time.

In recent years, governmental agencies, elected officials, emergency response agencies, and "first responders" have been diligently preparing for potential terrorist attacks. These preparation activities have been largely coordinated through the U.S. Department of Justice (DOJ) and the Federal Emergency Management Agency (FEMA). A particular focus of these efforts has been on protecting the governmental and business IT infrastructure.

In 1998, the U.S. Department of Justice Office for State and Local Domestic Preparedness Support convened the *Domestic Preparedness Stakeholders Forum* in the interest of assessing the level of preparedness to effectively and efficiently respond to a terrorist attack and to make recommendations regarding enhancing our level of preparation and readiness. This multidisciplinary panel was constituted to include representation from all of the agencies and major organizational entities that would be involved in responding to a major terrorist attack, particularly one involving weapons of mass destruction (WMD).

This panel identified a number of areas of our domestic preparedness for acts of terrorism that should be addressed, including comprehensive planning, coordination, incident management, training, equipment, and the interoperability of communications systems. A series of recommendations with respect to these issues was developed and presented to Attorney General Janet Reno and FEMA Director James Lee Witt. Both of these key members of the administration and the agencies that they led embraced the ever-present reality of the threat of domestic and international terrorism and embarked on a mission of enhancing our level of national preparedness.

The diligence of these federal agencies and others, coupled with their state and local partners, resulted in the development of a comprehensive system of response preparedness. The key to the success of this system is that it involves a coordination of local, state, and federal agencies and resources. This is imperative in that the initial response to any emergency incident will consist of local fire, law enforcement, emergency medical services, and emergency management resources.

Through a coordinated system of preparation and response, governmental entities and emergency management agencies have been preparing for inevitable terrorist attacks. The early preparation activities focused on our nation's 120 largest metropolitan areas. Later these initiatives were expanded to encompass 151 cities and/or metropolitan areas. More recently, these efforts have been further expanded to address the preparation of jurisdictions throughout the nation.

These preparation initiatives are designed to "plan for the worst" but assume that "the worst will never happen." It is indeed fortunate that governmental and emergency management officials had planned for the worst given that the worst did occur on September 11, 2001.

The fact that New York City and Washington, D.C. have been aggressive in their preparedness initiatives was also crucial to the outcomes of the events of September 11, 2001. The horrific events of this tragic day reinforced the importance of effective risk management and emergency planning and preparedness. While the loss of life on September 11, 2001 was massive, the ability of "first responders," including fire, law enforcement, and emergency medical services personnel, to evacuate and rescue so many individuals from the World Trade Center buildings attests to the importance of planning and preparedness. It is likely that if the events of September 11, 2001 were to have occurred on February 26, 1993, the day of the earlier World Trade Center bombing, the loss of life would have been much greater.

The importance of governmental agencies assessing vulnerability to acts of terrorism is evident, as is the importance of their planning and preparation. While as a nation we have made great strides in this pilgrimage, it is imperative that contemporary organizations likewise assess their vulnerability to acts of terrorism and engage in effective risk-management practices to address these vulnerabilities. An integral aspect of these initiatives must address the vulnerability and preparedness of an organization's IT resources.

The focus of this article will be on how assessment of organizational vulnerability and organizational preparedness has become a priority of most contemporary organizations since September 11, 2001. After relating a number of dimensions that should be addressed in assessing organizational vulnerability to acts of terrorism, we will examine what contemporary organizations and their IT functions can do to decrease their vulnerability to acts of terrorism and increase their level of preparedness to prevent, respond to, and recover from terrorist attacks.

Reality of the Terrorist Threat

The events of September 11, 2001, have clearly illuminated the fact that we as a nation, like other civilized countries around the world, face an ever-present threat of terrorism. Only through comprehensive and coordinated planning and preparedness can nations, communities, and organizations successfully prepare to defend against and respond to acts of terrorism.

Potential Threats

The five potential threats of terrorism that have been utilized in planning models have traditionally included biological events, chemical events, explosive events, incendiary events, and nuclear events. The potential for cyberterrorism has been recognized in recent years as a formidable means of terrorist attack. While some of these types of attack are certainly more probable than others, contemporary organizations should consider all six (including cyberterrorism) as they assess their vulnerability and preparedness.

Potential Harm

The harm resulting from a terrorist attack, invoked by one of the aforementioned means, can include asphyxiation harm, chemical harm, etiologic harm, mechanical harm, radiological harm, and thermal harm. These six types of harm are based on the five traditional types of terrorist attacks referenced and focus on harm invoked in terms of injuries, fatalities, and damage to tangible resources, including buildings, equipment, inventory, and infrastructure. The addition of cyberterrorism to the previous list should initiate contemplation of additional harm that an organization could suffer in terms of both tangible and intangible IT resources and infrastructure.

Vulnerability to Acts of Terrorism

Organizational vulnerability to acts of terrorism is influenced by a number of factors including: potential targets, environmental factors, organizational exposure, organizational preparedness, and business impact. Each of these dimensions must be considered as the contemporary organization engages in an assessment of its vulnerability to and preparedness for potential acts of terrorism.

Potential Targets

The potential targets of a terrorist attack can be categorized as general targets, organizational targets, and facility targets. In assessing their vulnerability to acts of terrorism, contemporary organizations must consider their organization in light of each of these types of potential targets. It must be understood that an organization may be a target from one or more of these perspectives.

General targets include controversial businesses, historical sites, infrastructure systems, places of assembly, public buildings, and symbolic targets. The World Trade Center and the Pentagon each involved a number of these characteristics of a general target. *Organizational targets* are based on the mission, products, services, political affiliation, and/or sociocultural position(s) of an organization. A given organizational facility may be a *facility target* based on its geographic location, proximity to other targets, products produced, services delivered, potential impact, or visibility.

Environmental Factors

A number of environmental factors may account for the ever-present reality of terrorism in general and for the vulnerability of specific organizations and their facilities to acts of domestic or international terrorism. These factors fall into the known general environment categories of political/legal, sociocultural, economic, and technological. The importance of the technological dimension must be emphasized given its relevance to the infrastructure utilized by and the business operations conducted by contemporary organizations.

Organizational Exposure

The two major determinants in assessing organizational exposure to vulnerability to acts of terrorism are the degree of exposure and the impact of exposure. *Degree of exposure* is influenced by mission-critical activities, centralization/decentralization, involved operations, and task interdependencies. In evaluating the *impact of exposure*, an organization should evaluate its ability to operate after a terrorist attack, its ability to survive after an attack, increased post-attack operating costs, and potential length of post-attack business interruption.

Organizational Preparedness

In evaluating organizational vulnerability to acts of terrorism, the current level of preparation to prevent, respond to, and recover from a terrorist attack must be fully understood. While an outcome of the process of assessing organizational vulnerability to acts of terrorism will likely be the development and implementation of organizational strategies to enhance organizational preparedness, the current level of organizational preparedness must be a key consideration in the evaluation of organizational vulnerability. While any number of attributes of the organization could prove to be either an organizational strength or an organizational weakness with respect to organizational preparedness, the areas of organizational planning, organizational structure, organizational controls, and risk-management programs are typically considered.

Organizational planning processes that provide for both comprehensive and contingency planning are desirable to enhance organizational preparedness, whereas ineffective or nonexistent planning compromises organizational preparedness. A decentralized *organizational structure* is an integral attribute of organizational preparedness, whereas a centralized organization contributes to an enhanced vulnerability. *Effective organizational controls*, in addition to ensuring an effective and efficient business operation, serve as a crucial component of organizational preparedness, thus organizations with ineffective control processes face an increased vulnerability.

Last, but certainly not least, are the *risk-management programs* currently operational within the organization. Effective risk-management programs are central to enhancing organizational preparedness for acts of terrorism, whereas ineffective risk-management programs or the absence of these programs greatly reduce organizational preparedness and increase organizational vulnerability.

Business Impact

In addition to the aforementioned considerations, the potential business impact of a terrorist attack must be considered in the evaluation of organizational vulnerability to acts of terrorism. This is a particularly important set of factors to consider as the contemporary organization seeks to realistically assess its vulnerability and to engage in proactive planning and preparation initiatives designed to enable the organization to prevent, respond to, and recover from a terrorist attack.

Factors to be considered with respect to business impact include business failure, customer/contract loss, personnel recruitment and retention, reduced market-share, reduced sales/revenues, reduced profitability, and reputation loss. One or more of these issues are likely to prove to be formidable issues for most contemporary organizations.

Rethinking Organizational Vulnerability and Preparedness

A recent research study conducted by the author sought to determine if contemporary organizations are rethinking their vulnerability to acts of terrorism in light of the recent tragic events. Corporate executives participating in the survey were asked to indicate if their perceptions of organizational vulnerability to acts of terrorism had increased, remained unchanged, or decreased since the tragic events of September 11, 2001.

Organizational Vulnerability

The majority of the survey respondents (87.3 percent) indicated that they recognized an increased organizational vulnerability to acts of terrorism since the events of September 11, 2001. The remaining 12.7 percent of the respondents indicated that the vulnerability of their organization had not changed based on the events of September 11, 2001.

Survey participants were asked if they expected the vulnerability of their organizations to acts of terrorism to increase, remain constant, or decrease over the next three years. The majority (89.8 percent) indicated that they anticipated that the vulnerability of their organization would increase over the next three years, while 10.2 percent indicated that it would remain constant.

Organizational Preparedness

The corporate executives participating in the research study were asked to indicate whether their perceptions of organizational preparedness for acts of terrorism had increased, remained unchanged, or decreased since September 11, 2001. The majority of survey respondents (82.2 percent) indicated that they perceived a decreased level of organizational preparedness, while 17.8 percent indicated that their perceptions of organizational preparedness had not changed since September 11, 2001.

Survey participants were asked if they expected the preparedness of their organization for acts of terrorism to increase, remain unchanged, or decrease over the next three years. A majority of the respondents (93.2 percent) expected that organizational preparedness would increase and only 6.8 percent anticipated that it would remain unchanged.

Changing Paradigms for Contemporary Organizations

The events that occurred on September 11, 2001, have since caused us as individuals, communities, organizations, and a nation to rethink how we live, work, and travel. Contemporary organizations are facing the reality of a new paradigm of conducting business. In recent years, we have heard highly respected and successful business leaders and academicians advocate that we must change the paradigms of how we do business. Our businesses have embraced the importance of quality management and made a commitment to continuous improvement and our economy has demonstrated the results. As with anything that is new, the quality movement had its early adopters, those organizations who bought in later, and those who held out and become late adopters.

The tragic events of September 11, 2001, served as the catalyst for a less subtle and much more timely paradigm shift. Contemporary organizations around the world are rethinking the way they conduct business. They are assessing their vulnerability to acts of terrorism and developing and implementing strategies to enhance their organizational preparedness for the formidable enemy of domestic and/or international terrorism. While for many organizations these analysis and planning processes are currently in process, it is anticipated that the outcomes of these proactive initiatives will include changing paradigms with respect to both business practices and conduct and the contemporary IT function.

Business Practices and Conduct

The ever-present threat of domestic and international terrorism, while a challenge that must be addressed by the leaders of contemporary organizations and the organizations that they lead, is not excepted to completely change the way in which the world of contemporary business functions or the business practices and conduct of a given contemporary organization. A proactive approach to embracing the threat of terrorism faced by the organization will, however, cause it to rethink a number of issues regarding the conduct of its business.

Organizational Mission: The mission of an organization drives organizational outputs, processes, and inputs. Organizational mission may cause an organization to become an organizational target. Contemporary organizations are thus advised to re-examine their mission in light of the threat of terrorism and revise their mission statement as appropriate.

Products and Services: The products produced and/or services delivered by an organization may make it an organizational target or facility target. A cost-benefit analysis may reveal that certain products and/or services currently offered contribute to an unacceptable vulnerability to terrorism.

Diversification: Contemporary organizations engage in a variety of diversification initiatives, including horizontal, vertical, and global diversification. They accomplish diversification through internal development, acquisition, and strategic partnerships. Contemporary organizations, interested in reducing their vulnerability to acts of terrorism, may consider making changes in their portfolio of business holdings and operations.

Organizational Planning: An integral component of the successful contemporary organization is its planning processes. The ability to manage in a proactive manner is based upon having comprehensive planning processes in place. The importance of assessing vulnerability to and preparedness for acts of terrorism while analyzing organizational strengths and weaknesses in light of environmental opportunities and threats cannot be overstated. Contingency planning must be addressed within the organization's planning processes.

Organizational Structure: The degree of centralization/decentralization of an organization can be a determinant of its ability to prevent, respond to, and recover from a terrorist attack. The more centralized the organization, the more vulnerable it will be to acts of terrorism. The more decentralized the organization, the less vulnerable it will be. The degree of centralization/decentralization can thus influence the various factors previously described with respect to the business impact of a terrorist attack. Organizations concerned about reducing their vulnerability are likely to implement strategies designed to further decentralize their operations and to enhance redundancy in the interest of business continuation should the organization experience a terrorist attack or other major emergency situation.

Organizational Controls: Organizations that lack control systems or have ineffective controls are likely to be more vulnerable to acts of terrorism than those that have effective control systems. It is imperative that contemporary organizations examine their control systems in light of the potential threat of terrorism.

Risk-Management Programs: The importance of having an effective risk-management program in place to protect the organization from the threat of terrorism should be obvious. This risk-management program should include both risk-assessment and risk-control measures.

Facility Site Selection: The physical location of a facility or its proximity to other targets can make one or more facilities of an organization a facility target. Potential vulnerability to acts of terrorism and the ability of an organization to defend itself against such attacks should be a new consideration when making facility site selection decisions.

Personnel Recruitment and Retention: The success of any contemporary organization is largely determined by its ability to attract, motivate, and retain a qualified and committed workforce. This will present an increased challenge given the concerns of existing and potential employees with respect to terrorism. Contemporary organizations must make the safety and protection of their personnel a priority.

Increased Use of Telecommunications: Concerns regarding terrorism have greatly reduced travel since September 11, 2001. Organizations are limiting traditional business travel and conducting more business meetings through the use of telecommunications. A benefit of this new approach to conducting business may be lower meeting costs and greater productivity. It is also expected that there will be a continuing trend toward the use of e-commerce.

The Contemporary IT Function

Each of the issues related to business practices and conduct just presented may have an impact on the organization's IT function. The decision-making involved in addressing these business issues will require IT support. Organizational decisions to change the mission, structure, products and/or services, or controls of the organization or to diversify its operations will pose many challenges for the IT function. IT resources will be central to the organization's risk management, human-resources management, and site-selection programs. The increased use of telecommunications will also result in increased stakeholder expectations regarding the organization's IT function.

In addition to responding to the IT challenges related to the changing business paradigms just discussed, the IT function within the contemporary organization will need to address the following issues regarding the provision of IT services that are effective, efficient, responsive, and secure.

Changing IT Role: The role of the IT function within the successful contemporary organization has undergone a significant evolution in recent years. The role and importance of the IT function are likely to be enhanced as a result of the new vulnerabilities that contemporary organizations are recognizing with respect to organizational vulnerability to terrorist attacks, including those involving cyberterrorism.

IT Resources That Support Business Needs: The IT function within the contemporary organization will be expected to continue to support the business needs of the organization. The importance of protecting organizational IT resources from attack will be paramount in both the present and the future. The preparedness of a contemporary organization to prevent, respond to, and recover from a terrorist attack will largely depend upon the security, redundancy, and robustness of its IT resources.

Recruitment and Retention of IT Personnel: The recruitment and retention of IT personnel is likely to become even more of a challenge in the future. Comprehensive screening approaches will be necessary to insure the preservation of IT resources through effective hiring and security practices.

Enhanced System Security: System security will become an increased priority in contemporary organizations. Vulnerability to traditional acts of terrorism, coupled with the potential of cyberterrorism, will result in the invocation of proactive strategies to ensure system security and the protection of organizational IT resources.

Protecting IT Resources: The importance of protecting IT resources should be obvious in light of the recent tragic events. The IT function will be expected to demonstrate proper stewardship for the organizational IT resources that have been entrusted to it. Strategies must be developed that ensure the protection of all IT resources, including hardware, software, and data.

System Configuration: Those responsible for the organization's IT resources should rethink the current configuration of their systems. This may prove to be a difficult and problematic issue. The traditional control and security advantages associated with centralized systems are now offset by the major disadvantages of IT resources that are primarily in one central location. Clearly, those responsible for protecting the physical systems in place within the World Trade Center buildings would likely not have imagined the physical attack experienced on September 11, 2001. Decentralized and distributed systems, which are properly secured, will likely be advantageous as an organization seeks to develop its preparedness to prevent, respond to, and recover from a terrorist attack.

Information System Controls: As always, information system controls must be our first line of defense in protecting IT resources. These control systems should include general controls, application controls, and network protection. The traditional types of general controls, including physical controls, access controls, data-security controls, communication controls, and administrative controls, will increase in importance. Application controls, designed to protect the system during input, processing, and output activities, will continue to be a mainstay of IT resource protection. Last, but certainly not least in light of the threat of cyberterrorism, will be the enhanced use of network-security measures and firewalls.

Disaster-Recovery Planning: Disaster-recovery planning will be central to an organization's preparedness for acts of terrorism. The ability of the organization to "get back in business" will determine the impact of its exposure to a terrorist attack. Organizations that are capable of restoring their mission-critical information systems in a timely manner will be able to restart their operations, survive, minimize increased operating costs, and limit the length of business interruption. This ability will consequentially reduce the potential business impact of a terrorist attack, including business failure, customer/contract loss, personnel recruitment and retention, reduced market-share, reduced profitability, reduced sales/revenues, and reputation loss.

Back-Up Systems: It goes without saying that the recent tragic events clearly illustrate the importance of system back-up routines that are religiously practiced, as well as the importance of off-site storage of back-ups.

Auditing Information Systems: Information systems should be audited periodically to ensure that they are operating properly in accordance with the business needs of the organization and acceptable business and information-processing practices. The issues considered in these audits must be expanded to include those practices that will reduce organizational vulnerability to acts of terrorism and will enable an organization to be prepared to prevent, respond to, and recover from a terrorist attack.

Preparing for a Successful and Safe Organizational Future

Contemporary organizations face many unprecedented challenges as a result of the ever-present reality of terrorism. The future success of the contemporary organization, the safety of its personnel, and the protection of its tangible and intangible assets will require a proactive management approach. Assessing organizational vulnerability to acts of terrorism and an organization's level of preparedness to prevent, respond to, and recover from a terrorist act is not an option.

Within the contemporary organization, we are all in this together. Corporate managers and the employees of the organizations they lead must embrace and aggressively respond to this challenge together. IT professionals must stand shoulder to shoulder with the organization's management team to ensure our readiness and defense. We stand together in the war on domestic and international terrorism.

Telecommunications Regulation in Puerto Rico: Past, Present, and Future

Phoebe Forsythe-Isales
President
Telecommunications Regulatory Board of Puerto Rico

The history of the telecommunications field in the Commonwealth of Puerto Rico is inseparably tied to that of our brethren in the continental United States. There is a shared monopolistic past; a common submission to the rule of law, be it federal or state; and a strong sense of fairness, all of which have given rise to corresponding approaches to the regulation and deregulation of telecommunications markets.

In 1934, during the Great Depression, Congress enacted the Federal Communications Act. The 1934 Act established an independent regulatory body, the Federal Communications Commission (FCC), with broad powers to make nationwide wire and radio communications services available to all at reasonable prices and through adequate facilities. However, the 1934 Act's scope was limited to interstate and foreign communications, leaving intrastate regulation in the hands of local agencies, usually called public service or public utility commissions (PSCs or PUCs).

During the decades that followed the adoption of the 1934 Act, the intrastate telecommunications market was assumed, both by federal and state regulators, to constitute a so-called natural monopoly, that is, a market that can best be serviced by a single local-exchange carrier (LEC). Under this view, the role of the state PUC was, fundamentally, to outlaw competition and authorize the incumbent LEC (ILEC) to charge high business rates while simultaneously pushing said incumbent to maximize telephone penetration in low-cost and often unprofitable residential markets. This meant that business customers subsidized residential customers.

In Puerto Rico, the telecommunications market had displayed its monopolistic colors since its origins in 1914. The high cost associated with the establishment of infrastructure as well as other technical difficulties left a deep monopolistic imprimatur. The monopoly eventually extended to local service and intra-island long-distance service and became government owned in 1974. By the same token, AT&T's stateside monopoly not only integrated local and long-distance service, but telephone manufacturing as well.

Excluding the federal antitrust action against AT&T during the late 1970s, which in 1984 brought about the breakup of AT&T from its regional Bell operating companies (RBOCs), local monopolies remained generally undisturbed until the landmark passage of the Federal Telecommunications Act of 1996.[1] The 1996 Act radically shifted the telecommunications paradigm that structured local telephony. Not only did the 1996 Act outlaw local barriers to market entry and competition, but ILECs would now bear substantial obligations toward competitors. Under the new regime, ILECs would now have to allow rivals to interconnect with its network elements (switches, loops, trunks, etc.) on an unbundled basis and at advantageous cost-based rates. ILECs would also be required to facilitate market entry by new competitive LECs (CLECs) by providing them, at wholesale discount rates, services that it normally sold retail to its subscribers. Upon a finding that RBOC ILECs had taken meaningful and sufficient steps to open local markets, the FCC would then allow ILECs to enter into the manufacturing and interstate long-distance markets.

The 1996 Act enabled PUCs and PSCs in the 50 states and Puerto Rico to steer the implementation of the newly deregulated intrastate telecommunications landscape. The local PSC would make sure that the ILEC abided by the new cost-based interconnection and wholesale service rules, as well as enforce other market-opening powers, such as local number portability (LNP) and dialing parity requirements, prohibition of the *cramming* of unwanted services into a user's package, proscription of the illicit change of a user's presubscribed carrier (*slamming*), regulation of rights of way and collocation, mediation and arbitration of interconnection agreements, reciprocal compensation arrangements, among others.

In Puerto Rico, the monopoly of intra-island service persisted. As a result, the Commonwealth Legislature saw fit to create a specialized and dedicated body to carry out the local aspects of the 1996 Act's mandate. Unlike stateside PUCs and PSCs, which oversee several utilities such as electricity, gas, water, and telecommunications, Puerto Rico opted for an agency established for the sole purpose of achieving the pro-competitive goals set by Congress, while minimizing the hindrance and complication that involves regulating multiple and radically different sectors of the economy. As such, and just seven months after the Federal 1996 Act was passed, the Governor of Puerto Rico signed into effect Act 213 of September 12, 1996, otherwise known as the Puerto Rico Telecommunications Act of 1996, the statute that created the Puerto Rico Telecommunications Regulatory Board (the Board).[2]

The Board was created with the powers and prerogatives ceded by the FCC to PUCs as per the 1996 Act and with a clear mandate to foster total, equal, and fair competition. As a direct consequence of the 1996 Act and of Act 213, the Commonwealth's incumbent, PRT, faced, for the first time in its 80-year history, the prospect of sharing its turf with new market-entrants. Specifically, the Board was constituted to a) ensure the availability of universal telecommunications services at affordable rates for all citizens of Puerto Rico; b) oversee the efficiency of telephone service and cable television, as well as other telecommunication services; c) assure the continuity in the rendering of services of a social nature, such as public telephones, according to the public need; d) promote competition; e) guarantee Puerto Ricans the same telecommunications and information privileges enjoyed by other United States citizens; and f) safeguard the public interest to the utmost. The Board, a three-member regulatory body, was not only instituted as a quasi-legislative entity, but also as a well-staffed, quasi-judicial, adversative and deliberative organism capable of receiving and resolving inter-company and consumer-to-company complaints.

Throughout the past five years, the Board has undertaken numerous initiatives in its quest to carry out its statutory duties. Amongst the most recent lies the commencement of telecommunications relay services (TRS) for citizens with speech and auditory impediments; the overlay of an additional area code (939) that shall ensure continued and sustainable growth in the Commonwealth's telecommunications industry (an area code equals 7,920,000 new telephone numbers); and the implementation of dialing parity or local-toll presubscription, which gives consumers the same options when selecting a telecommunications provider for local toll calls as when choosing a long-distance carrier (interexchange carriers [IXCs]). Furthermore, the Board, in conjunction with the Puerto Rico Education Department, has procured $115 million in federal Universal Service *School and Library Program* funds for the installation of facilities and connection of 1,540 public schools to the Internet and the World Wide Web (WWW) and has engaged in an island-wide promotion of federal Universal Service *Low Income* programs, particularly *lifeline* and *link-u*p, which provide telephony subsidies to people of lesser means who are on welfare.

The Board is making every effort to move Puerto Rico closer to the most economical and forward-thinking U.S. jurisdictions when it comes to the provision of telecommunications services. The Board strives to utilize competition and market forces as the primary setters of the prices, terms, availability, and conditions of telecommunications services, as long as prices are based on cost. Companies should always seek negotiated solutions for controversies among them. Conversely, decisive Board intervention is often required to rectify market iniquities. Cross-subsidization between competitive and noncompetitive services, price schemes that artificially sustain unreasonably low rates, and myriad other vestigial monopolistic and anti-competitive practices must be categorically eliminated in order to reap the full benefits of market liberalization. The Board's rule-making and adjudicative endeavors are slowly crafting a fair playing field that promotes competition and cost-based pricing, so that consumers pay for the services they are actually receiving.[3]

The aforementioned measures are, however, but first steps toward the realization of a much greater and ambitious goal spearheaded by the Board: the transformation of Puerto Rico into a hemispheric telecommunications hub, principally for the benefit of our Latin American and Caribbean environs. Puerto Rico possesses a 100% digital network, lies under the footprint of 92 satellites[4], is a landing spot for six major marine cables[5], and enjoys a substantial deployment of fiber-optic rings. The Board has been honored by the Caribbean Association of National Telecommunication Organizations' (CANTO)[6] acknowledgment of its regional relevance and the president of the Board has gladly provided technical assistance to telecommunications commissioners from the Dominican Republic, Jamaica, the Republic of Honduras, and the Commonwealth of the Bahamas, among others.

All in all, the Commonwealth's pro-competitive mindset, its comparatively advanced infrastructure, its stable government and financial institutions, its advantageous geographic location, its close relationship to the United States, and its cultural ties to the Spanish-speaking world provide unique opportunities for new entrepreneurs and already-established companies.

The Board is hard at work, setting in motion processes that are raising Puerto Rico's attractiveness as a home base for hemispheric telecommunications operations. Under the Board's auspices, Puerto Rico has witnessed an impressive expansion in the availability and quality of telecommunication goods and services. Two hundred and thirty three telecommunications companies have become established in the island during the last five-year period, and Puerto Rico boasts the highest per-capita Internet penetration rate in Latin America and the Caribbean. Competitive markets tend to drive down rates 40% to 50% lower than noncompetitive monopolies. As such, rates should continue their island-wide decrease across the whole spectrum of telecommunications services and particularly in the wireless market, where price wars benefiting the consumer are the order of the day. Although the ideal of total, equal, and fair competition has not yet been achieved, it can be said that the era of legally justified natural monopolies has come to a close in Puerto Rico.

The future of the telecommunications industry, which already is a $600 billion market, is quite bright. According to the FCC's latest numbers, CLEC market-share in the United States, including Puerto Rico, grew 93% over the one-year period of January to December of 2000, and the number of wireless service subscribers was 101 million, up from the previous 91 million reported for 1999. At least one CLEC was serving customers in 56% of the nation's zip codes, and about 88% of U.S. households reside in these zip codes.[7]

As per industry estimates, Puerto Rico boasts close to 1.7 million telephone lines, approximately 1,500,000 wireless phones in use, 551,600 Internet users, and 374,000 cable-television subscribers. The wireless market alone has experienced close to 300% growth during the past five years. In Puerto Rico, 54% of families own wireless phones, at around 1.6 units per home. These numbers will certainly rise, specifically with the deployment of nascent technologies such as digital subscriber line (DSL) services, cable modems, and third-generation (3G) wireless systems.[8] There is also ample space for growth in landline service. Island penetration lies at around 74.6%—quite lower than the highest level experienced stateside, which is actually 94.5%.

The new paradigm brought forth nationally by the 1996 Act and locally by Act 213 is already delivering substantial qualitative and quantitative movement toward a future in which telecommunications companies spur technological progress and ever-more complex services while battling for customers in a highly competitive marketplace. The final result is gain for all: Companies grow and profit, and Puerto Rico benefits from industrial expansion and higher employment. Meanwhile, the consumer receives high-quality products at excellent prices and has myriad companies from which to choose. The Board, on its part, shall continue to guard the legal regime that has made all of the foregoing possible.

Notes

1. Telecommunications Act of 1996, P.L. No. 104-104, 110 Stat. 56; 47 USC §§ 151 *et. seq*.
2. Act 213 of September 12, 1996, 27 L.P.R.A. § 265 *et seq*.
3. For general information or for further detail regarding ongoing legal proceedings visit the Board's Web site at http://www.jrtpr.gobierno.pr.
4. Among these satellites are 1) Aurora, 2) Morelos, 3) Satcom, 4) Solidaridad, 4) Galaxy, 5) Brazilsat, 6) Tel Star, 7) Intelsat, 8) Echo Star, and 9) Thor.
5. These marine cables are 1) Americas, 2) Arcos, 3) Florico II/TCSI, 4)Taíno, 5) Antillas, and 6) SAM USA.
6. CANTO is a regional telecommunications organization based in the Republic of Trinidad and Tobago that consists of 47 members, generally national operating telecommunications companies from the Caribbean and South, Central, and North America. Membership is open to operating companies in the region. A board of directors drawn from the membership governs the organization.
7. *Federal Communications Commission Releases Latest Data on Local Telephone Competition*, http://www.fcc.gov/ccb/stats.
8. Promising 3G technologies can transmit text and digitized audio and video at speeds up to 2 megabits per second (Mbps) in a stationary wireless service and at 384 kilobits per second (kbps) in a mobile environment. This allows for services such as fast wireless Internet and real-time videoconferencing.

DSL: The Customer Perspective

Duncan Greatwood
Vice President, Marketing
Virata Corporation

The telecommunications industry sometimes seems obsessed with improving the business case for digital subscriber line (DSL). In many ways that is a good mindset, as the industry will grow because of a strong business orientation. That approach, however, tends to ignore the other side of the story—the customer perspective on DSL. This paper looks at DSL from the customer side and examines the aspects of DSL that customers find especially appealing.

Cost Reduction

First of all, DSL customers are looking for lower cost—both lower equipment costs and lower service-delivery costs.

Reducing Equipment Costs

Today's equipment already contains local-area network (LAN) interfaces. The DSL transceiver, analog chip, and discreet circuitry are built in, enabling a cost reduction through integration. Digital integration is fairly well advanced, as today's equipment already provides the necessary level of integration of the transceiver, the processor, and the LAN interface. However, analog integration is more challenging because it does not scale down as easily into the silicon geometry as digital does. At this point, maintaining performance becomes an issue.

One solution is to optimize analog circuitry. The analog discreet circuit involves substantial cost, and analog optimization can provide substantial savings. The circuit can be integrated into the analog chip, including line drivers, receivers, and the other components. In addition, some of the functions implemented in analog circuits can be transferred into digital formats. When these analog functions are dropped into the digital chip and integrated with the transceiver, the solution is more cost effective because digital silicon costs less.

Thus if the software and hardware implementations are correctly partitioned on analog and digital chips, the best and highest-performance solution can be found. This solution is a system-wide one—and building the best system as a whole is the issue. Using hybrid technology to do so is much less expensive than a full-system design using a single protocol such as Internet protocol (IP).

Service-Delivery Costs

In the early days of DSL, service-delivery costs were extremely high. The costs included a truck roll to install the service, and one truck roll was seldom sufficient to get the copper wiring to work correctly. In addition, interoperability issues emerged between different DSL systems.

Those problems have, for the most part, been solved. DSL installation now means that the user installs pre-configured software and filters. The result is a more consistent user experience, a polite way of saying that DSL now works as expected on a regular basis. Interoperability in the physical layer has also removed some of the barriers that existed just a few years ago.

Volume continues to grow, however, and with it an important cost for service providers: the huge cost of digital subscriber line access multiplexers (DSLAMs). Installing the DSLAM to provide service generally involves a cash-flow problem for providers, who must invest a great deal at the front end and do not begin to recoup their investment until the service is deployed. One approach to this problem is to reduce the time gap between the investment and the service.

Thus, although service costs are declining, more remains to be accomplished.

Increasing Customer Value

While customers are concerned with costs; their perception of the value of the service delivered must also be considered. Value can develop in two ways. First, the service provider can provide a service directly and charge for it. For example, the provider can offer voice-over–DSL by providing multiple lines of voice over the DSL wire, pulling them off over the DSLAM, and routing them. This approach results in actual revenue streams. Another way to produce value is by providing applications that do not require the DSL service provider to participate at all. Examples of this type of application include Internet games or music download services. There is a whole category of broadband applications that flow over the DSL infrastructure but do not require the DSL service provider to take any action. This category is clearly of value to the service provider. Instead of paying $30 a month, consumers may be willing to pay $35 or even $50 for the extra value they perceive, and that makes a big difference in the provider's bottom line.

Building value in the eyes of the consumer is a significant change for service providers. DSL acceptance will accelerate as customers perceive it not only as a way to download the same things faster, but also as a way to do things that they otherwise could not. Building value in this way offers a much more compelling proposition.

Possible Applications

Anybody can deliver service. What will differentiate DSL providers is generating revenue in both the ways outlined above—by offering services themselves and by enabling other services that do not require their active participation. Voice, as briefly mentioned, is one such application, but because it has been available for so long it is no longer as compelling as some of the newer applications. New uses will drive demand more quickly. Broadband games will generate huge demand; home security is another big area, as the 6 million physical-security units sold annually could be easily converted to DSL. Photography storage and download capabilities are other areas that add value to DSL from the customer perspective.

The Future

DSL promises to transform the home through expanding uses of broadband technology. The broadband home will utilize multiple devices to access the Internet. This scenario will require a very simple, very easy-to-use network infrastructure in the home. Everything from the computer printer to the camera to the stereo system will be connected to this network to access Internet applications, thereby driving DSL usage.

To make this vision possible, customer-premises equipment (CPE) must become a platform for all the home uses envisioned. Today, CPE is regarded as a commodity, and providers have done well selling it as such. Increasingly, customers will want a platform they can use to build their own value-added services. DSL will support this view by working consistently and effectively with all other parts of the system at a low cost. Value creation will occur at the application level, in the digital loop carrier (DLC,) and in the DSLAM.

Among other requirements for enabling the application requirements, modularity will be especially important. Services such as asynchronous transfer mode (ATM) and IP for communications processing will have to be available through a range of different types of consumer equipment. Eventually, customers will be able to choose from a wide variety of equipment for different needs in a basic retail model.

Conclusions

DSL providers that consider the customer perspective will offer a better value proposition. The customer perspective involves reducing costs. Equipment costs can be reduced by providing an integrated hybrid system. Service-delivery costs are already declining, but more remains to be done to truly meet customer needs. Adding value through service delivery is the key to the customer perspective, and providers can add value by providing applications themselves, enabling other applications, or both.

When providers implement the customer perspective through lowering costs and delivering value-added services, they will provide the necessary support for using a CPE platform to make the broadband home a reality.

Choosing the Best Internet Data Center: Learning from the Japan Experience

Mike Heller
Senior Manager
Cisco Systems, Inc., USA

Kiochi Tamura
Manager
Cisco Japan KK, Tokyo, Japan

Introduction

Businesses of every size in countries all over the world are facing more challenges than ever before. They are finding increased competition and a faster pace of change. Their customers expect them to be open or accessible for longer hours (24 hours a day for many), to have the latest products and services more quickly, and always have competitive prices. These phenomena are present in Japan, just as in the United States, Europe, and other areas of the world.

The Internet, and in particular the World Wide Web (WWW), presents an opportunity for businesses to address these needs. An effective Web presence can help them to extend their reach and attract new customers, while lowering their costs of doing business.

When an organization investigates establishing a Web site, they quickly discover that this is not as simple as one might hope. An environmentally managed and secure room with a raised floor will be needed for the computer servers needed, along with an increased bandwidth connection to the Internet (dial-up connections are entirely unsuitable). But perhaps the most difficult challenge is that of hiring and retaining adequately knowledgeable staff. (Just to learn all the jargon and acronyms is a huge task!)

To help businesses take advantage of the Web, yet avoid these problems, a new industry has sprung up in recent years, sometimes called Web hosting and here referred to as Internet data centers (IDCs). IDCs may be operated by traditional telecommunications carriers or by specialist companies offering only those services.

IDCs typically offer a secure physical environment, perhaps a dedicated building or perhaps a secure area within a larger building. Twenty-four–hour staffed entry doors, key cards, or biometric personnel identification systems; video surveillance; and strict guest logging procedures are common. The data center will have uninterruptible power supplies (UPSs), along with redundant air conditioning and humidity-control systems. There will be a 24-hour network operations center (NOC), with a good database of the installed network and customers' systems, automated trouble ticketing, and escalation systems, etc.

If Your Business Relies on the Web, Your IDC Must Be Reliable

Using an IDC is not without some problems and concerns, however. Unless chosen carefully, the reliability and security of the IDC may be no better than the business could put together on their own, and this may be hard to ascertain. Some smaller IDCs may not be staffed 24 hours a day, making quick access to your servers difficult.

There are many IDCs in operation, so how does a company pick the right one? There are checklists available to aid in the investigation and comparison of the physical facilities of different IDCs that a firm may consider, and some key points are subsequently summarized. This paper complements those resources and adds to the chance of a more informed selection by allowing a company to compare the offerings of several IDCs through case studies.

The various checklists focus on two different categories of issues that can lead to greater reliability in an IDC. These are 1) the way in which the IDC has been designed and 2) the way in which the IDC is staffed and operated.

Be sure to visit any IDC being considered. Notice how easy it is for you to enter the building, what identification is required, whether you needed an escort, if you were constantly in view of video surveillance systems, and what records of your visit are kept. Ask about how electric power is supplied to the building and about what generators and other back-up systems are in place (and how much capacity they have). Redundant air-conditioning systems are a must in warmer climates (computers and other networking equipment gets rather hot), and check an effective fire detection and suppression system. What does the building look like? Would it likely withstand an earthquake, a tornado, or a flood?

An explanation of the way that data networks are designed should include how many different carriers connect to the building, how redundancy has been included in the network design within the building (there should be *no* single points of failure), and how many different equipment vendors are used (the more there are, the more complexity there is in operations, training, etc., and the greater the likelihood of problems, "finger-pointing," etc.).

Having said all that, more than 80 percent of telecom network outages are caused by operations errors or other "people problems," so be sure to assess how network operations is handled. An effective NOC will have clear and documented procedures for record keeping, trouble tickets, escalation procedures, exception reporting, working on customer equipment, etc. Successful management of an IDC requires careful attention to *configuration* management (e.g., software-version control, test plans, password management), *performance* management (e.g., capacity planning, quality-of-service [QoS] management, fault detection and resolution procedures, help-desk metrics, etc.), and *security* management (e.g., access control, partner and contractor management, intrusion reaction, etc.).

Management of an IDC means much more than just knowing if some piece of equipment is "up or down." The ability of the IDC staff to measure and assess service levels on a per-customer basis is essential to achieve the promised service-level agreements (SLAs). Managing Internet protocol (IP) address assignments, keeping track of network configurations, and not losing site of trouble tickets and alarms often requires the use of a mechanized NOC record keeping and support system. Good manageability tools and qualified personnel supporting the infrastructure translate into lower operations costs (since time is not wasted trying to resolve indications from conflicting management systems) as well as higher customer satisfaction.

If the service provider has more than one IDC, customers can benefit by having data and content at multiple sites, thus providing faster response times and higher availability (due to the partial or complete redundancy of the data). But it also adds complications for the IDC operations staff, including systems and procedures for coordinating the different centers, distributing updates to the proper locations, synchronizing data at different sites, handling downed servers, coping with the added security risks of remote access, etc.

The case studies presented in this paper on different IDCs in Japan often mention that SLAs are available to customers. A SLA can be a powerful tool to reach a common understanding with an IDC about service levels, availability, security, escalation procedures, etc. Many SLAs include provisions for some sort of refund or penalty payment if agreed service levels are not met. But it is also worth remembering that an SLA is somewhat like a life insurance policy, that is, you should never want to collect on it. Most SLA payments wouldn't cover the lost revenue or loss of customer goodwill that a prolonged outage would cause.

When choosing an IDC, it is always useful to ask how the customer will know what sort of performance levels the IDC is achieving. Regular reports from the IDC are a good sign, since a service provider that tends to want to hide this kind of information usually doesn't have service levels of which they are proud.

The Japan Business Environment

Business firms in Japan are mainly either very large or very small. Most of the very large ones are well known internationally, such as Sony, NEC, Mitsubishi, Honda, Toyota, Nissan, etc. Of course, not well known internationally are the hundreds of thousands of small "mom and pop" businesses and stores, but they constitute a large part of the overall economy.

The telecommunications industry in Japan changed from the classical monopoly environment, with NTT as the only provider, to a fully competitive market, with close to 10,000 service providers licensed by 2001. NTT, with government ownership now dropped to less than 50 percent, has been divided into several specialist companies (e.g., NTT Communications, handling long-distance and international traffic; NTT DoCoMo, handling wireless services; NTT East, providing local telephone services in the eastern part of Japan; NTT-ME, set up to offer Internet and "mutimedia" services; etc.). Principal competitors to NTT are KDDI, which grew out of the old international telecom monopoly company, and Japan Telecom, a new carrier that initially built its facilities along the rights of way of one of its parent companies, Japan Railways.

Taken together, these two factors greatly influenced the development of the IDC business in Japan, since there are many smaller companies that may not have large information technology (IT) staffs as prospective customers and many new service providers eager to find revenues in meeting new service needs.

IDC Developments in Japan

The first IDC opened in Japan several years ago, and by mid-2001, almost 100 companies had entered into the IDC business. Only a few larger companies, however, operate most of the data centers. The major Japanese telecom carriers are, of course, IDC operators, but some large system integrators and some global IDC specialist companies are as well. Some leading IDC operators in Japan are shown in *Table 1*.

At the end of 2000, Goldman Sachs Investment Research estimated there was a total of 95,000 square meters (slightly more than 1 million square feet) of IDC floor space available in Japan—enough, at average configurations, to accommodate about 210,000 Web servers. The same study found that only several hundred servers were actually in use in these centers, thus leading to many IDC suppliers with unused capacity pursuing the small but growing number of hosting customers.

While a few IDC operators have only a single data center, two or three data centers is more typical, and the larger operators have additional centers under construction.

It is believed that a Japanese IDC, to be successful, must include four key factors. The IDC provider must first provide a secure hosting facility. The second, and related, factor is interconnection with the Internet backbone to provide

TABLE 1

Some Leading IDC Operators in Japan

Japan Telecom Companies	Global Telecom Companies
NTT	Genuity
KDDI	WorldCom/Digex
Japan Telecom	Cable & Wireless/Digital Island
Internet Initiative Japan	Equant/GlobalOne
Biglobe	Level 3
System Integrators	**Other IDC Operators**
Fujitsu	Toppan Printing Company
NEC	Dai Nihon Printing
Hitachi	ClayFish Venture Capital
Toshiba	Chubu Telecommunications
CRC	

high-speed and high-volume data connectivity, especially to sites outside of Japan. Third on the list of success factors is a network and operations design (e.g., redundancies) to provide highly reliable and fast response serving of data. The last factor is a range of value-added services aimed at helping the business user to be successful at e-business.

An organization to represent IDC operators and promote this business model was founded in October 2000. It had grown to more than 150 companies by mid-2001. Some founding members include NTT, Sony, Oracle, Cisco Japan, Fuji, Sun Microsystems, Cable & Wireless, and GlobalOne. Called iDC, the organization runs advertisements and conducts seminars to educate potential users on the advantages of outsourcing Web servers and other business applications.

Although there are some differences in pricing between one Japanese IDC and another, the operators have mainly chosen to compete on the services that they offer, how they are packaged, what options are available, etc. Typical IDC monthly prices in Japan in late 2001, for various service options, are outlined in *Table 2*.

Japan IDC Case Studies

In this section, we present case studies of six IDCs now doing business in Japan. (There are other IDC services, but those with the most information available about their offerings and early experiences were selected for this paper.)

1. Digital Island Japan

Digital Island Japan is a subsidiary of Digital Island, headquartered in San Francisco, California. Cable & Wireless acquired Digital Island during 2001, but the Digital Island brand continues to be used. Its name reflects the company's original founding in Hawaii several years ago. Digital Island has a dedicated fiber-optic network connecting 35 countries and has IDCs located in the United States, Europe, and Hong Kong, in addition to Japan. Digital Island has grown from its founding in 1996 to a global business today housing more than 2,550 servers.

Digital Island is a specialist company devoted to IDC–type services. While basic Web hosting and co-location services are available, a major focus is on services for high-volume,

TABLE 2

Service	Monthly Charge (Yen)	Approx. U.S.$ Equivalent
Rental of rack space (1/4 rack)*	50,000	420
Rental of rack space (full rack)*	175,000	1,450
Shared server	4,000	35
Dedicated server	40,000	350
Shared Ethernet** connectivity	15,000	125
Dedicated Ethernet connectivity	500,000	4,200

**Includes limited technical support, e.g., visual inspections or rebooting of customer's equipment*
***10 Mbps; faster connections (e.g., 100 Mbps Fast Ethernet) may optionally be available*

high-performance content-delivery applications and e-business. For example, one service, called Footprint‰, is a content-delivery technology aimed at streaming media and music delivery. The present network can deliver data simultaneously to 300,000 users and is being expanded even further. It is particularly targeted to the financial market (real-time market data), the entertainment industry (especially music), the publishing industry (news), and the software industry (software downloads).

Services currently offered are shown in *Table 3*. Pricing generally consists of a monthly charge that varies with the amount of equipment provided by Digital Island and the responsibilities taken on by their personnel, plus a "data charge" calculated on the number of gigabytes of data that were served to the network during the month. Negotiated prices apply to any custom services requested, and volume discounts are available.

Digital Island offers SLAs to their customers. Typical components include network availability (e.g., 99.98%), network packet loss (e.g., less than 1%), latency (e.g., 80 milliseconds or less for Asia), and server availability (e.g., 99.8%).

2. NTT-ME

NTT-ME is a part of the NTT group and is focused on business applications. NTT-ME is directly connected to the XePhion2000, a gigabit-class high-capacity backbone network. NTT-ME utilizes the highly reliable equipment and facilities made available through NTT's know-how to build communications sites and networks of the highest standards. Their staff (which numbers more than 18,000 IT professionals) has advanced technical knowledge and can address a wide variety of customer broadband data transmission and hosting needs. For example, customers may optionally sign up for additional security and firewall services or for load balancing across multiple servers. NTT-ME also has a wide range of data storage options for situations when the client's data exceeds the amount of disk storage included with the hosting service packages. NTT-ME also offers an unusual service for IDCs, namely the complete hosting and support of Oracle business applications. NTT-ME has a partnership and technology-sharing arrangement with InterNAP in the United States, which facilitates the handling of the business of multinational companies.

In addition to basic Internet connectivity options, NTT-ME also can provide Internet protocol–virtual private networks (IP–VPNs), especially for clients of their business application software hosting services.

Customer satisfaction is important to NTT-ME. In addition to comprehensive SLAs, NTT-ME has partnered with several Japanese insurance companies to offer an unusual option: "Digital Business Insurance." These policies can fully cover all business losses (including lost revenue) that might occur under catastrophic circumstances, whereas typical SLA "penalties" are much less, perhaps a refund of a portion of the charges paid. A customer-support center and NOC operates 24 hours a day, 7 days a week.

NTT-ME has evolved from a pure IDC service provider to one offering not only hosting and related services, but also application service provider (ASP) services, particularly those related to Oracle applications and groupware, for which NTT-ME is the only IDC–based provider in Japan. These applications include a complete integrated package of Oracle applications, providing enterprise-resource

TABLE 3

Service Currently Offered

Service Type	Service Name	Features	Remarks
Server Hosting	Trusted Host™	Basic hosting of customer's server	Digital Island provides network connectivity equipment
	Extended Host™	Includes load balancing between servers and "hands-on" service functions	On-line access to server performance data
	Total Host™	Includes server and management up to operating-system level	
	Private Host™	Floor space and power only; Fast Ethernet connection	User supplies servers, routers, etc.
Streaming Media	Footprint™	Network-edge multiple-site caching (e.g., at ISPs) of content	Used for static data
	Footprint™ On-Demand	Same; for video clips & music	All media player types OK
	Footprint™ Live	Same; for real-time images	Options include image encoding
Localization	TraceWare™	Identifies location of requester and serves localized content	By country or region

management (ERP), supply-chain management (SCM), and customer-relationship management (CRM) functionality. They also include the "Intra-Mart" groupware package for use within companies to promote improved efficiency and information sharing. NTT-ME can offer a comprehensive and complete service, one bringing great value to smaller firms that need to concentrate their efforts on their core business and outsource the operations and management of business applications as well as the housing of their servers.

Although NTT is "the" telephone company in Japan, the NTT-ME subsidiary was a bit late to the IDC market, and in late-2000 held only about a 4 percent market-share, according to a survey by *Nikkei Communications* magazine.

3. KDDI

KDDI is a fairly new company, formed in 2000 by the merger of KDD, the traditional Japanese international telephone carrier, and DDI, a competitive long-distance and cellular operator started in the 1990s. Toyota, one of the original DDI shareholders, still has a small shareholding. When the Japanese market moved to a fully competitive mode, KDD was able to enter domestic markets, and DDI was then able to offer international services. An extensive set of data and Internet services rounds out the portfolio offered in Japan. A joint venture with Singapore Telecom provides services to business customers throughout the Asia-Pacific region.

The IDC focus of KDDI is large-scale e-business, although KDDI emphasizes that even small office/home office (SOHO) accounts are welcome. Under the brand name "dot-square," KDDI draws upon years of experience as a telecom carrier operating numerous data centers for billing and other purposes. KDDI continues to expand, with new data centers added around Japan in 2000 and a large new IDC planned for launch in Tokyo during 2002. All of the KDDI data center's feature redundant power systems, over-engineering air-conditioning equipment, and 24-hour monitoring by the on-site security staff.

Accordingly, KDDI stresses the robustness and capacity of its backbone network and its connectivity to Internet providers in other countries. KDDI's backbone network reaches customers' local-area networks (LANs) through direct connection via leased lines, frame-relay circuits, or IP–VPN encrypted tunnels across the Internet. SLAs are available for all services, and comprehensive performance reports are regularly provided to all clients.

The KDDI dot-square service offers a wide selection of service options to meet the unique needs of every size of customer. For the smaller customers, a basic hosting service is available, using redundant servers and circuits within the data center. Medium-sized customers, including ASPs and Internet service providers (ISPs), can utilize dedicated servers and rack space for routers and other equipment. These racks can be open or equipped with locking doors. Enclosed cages are available for larger customers or those desiring a greater degree of security for their equipment. The exact offerings vary somewhat depending on the size of the IDC but are similar in all. On-site staff in each center can perform "lamp checks" or other simple tasks at the customer's request.

Different Internet connectivity options are available for each hosting service, providing direct connections not only to the highest-capacity Japanese backbone, but also to the Japanese interchange points (JPIX and NSPIXP2) and to major U.S. ISPs.

KDDI offers a variety of operations and professional services, from its own staff and through partners, including in other countries. This can range from routine tasks such as mounting tapes and coordinating the work of contractors to major project management, network engineering, and software development.

According to a late-2000 survey by *Nikkei Communications* magazine, KDDI led the IDC market in Japan with an approximate 10 percent market-share.

4. CRC Institute

The CRC Institute has a large IDC located at Otemachi, near Tokyo. CRC has a proven track record in outsourcing and providing professional services to its clients, and the IDC services are a natural complement. With more than 10 years of experience in operating data centers in Yokohama and Kobe, the CRC Institute has the expertise and well-trained staff that customers expect from one of the pioneers of computer outsourcing in Japan. Their motto is "Comprehensive Solutions and Full Services."

An example of the unique services of the CRC IDC is the Internet security outsourcing service (iSOS), whose need has become apparent with recent hacking events. The iSOS service includes an "unauthorized access detection" capability—a form of what's generally known in the industry as an intrusion detection system, or IDS. It also has a proactive element called "security inspection," which has CRC security experts check the customer system periodically for security risks and vulnerabilities, and recommend improvements. A managed firewall service (called the CRC "firewall outsourcing service") is available, along with a "virus-protection service," which provides users with the latest antiviral software and assistance.

Of course, the CRC IDC services include 24-hour operations monitoring and site security (including comprehensive video surveillance), with various options that can be selected depending on individual customer need. The IDC itself is the latest design, featuring many earthquake-resistant aspects developed from the Kobe earthquake experiences. Power is supplied to the IDC from multiple substations and serves the building in a double loop feeding system. The UPS and on-site generators provide constant power, even during momentary outages or extended periods (e.g., legal inspections). The building air-conditional systems are redundant as well, and the equipment rooms and under-floor areas are served by gas fire-extinguishing units for quick detection and initial extinguishing of a fire for minimal damage.

Customers generally supply their own servers, routers, and other equipment, taking advantage of the extensive site support and network (and server) monitoring capabilities of the CRC IDC. A wide choice of rack, cage, and private-room options are available. Network connectivity is assured to be of the highest quality and performance due to the fact that seven different domestic and international network carriers have brought their facilities into the Ohtemachi building.

The CRC Otemachi IDC is furnished with various conference rooms and other facilities that clients may use for meetings with their customers, presentations, training sessions, etc.

Summary and Conclusions

This paper has overviewed the many benefits that a company can obtain by using the services of an IDC to host, and perhaps manage, their Web servers and Internet connectivity. Due to budget considerations and the difficulties of hiring a suitably trained staff, for many smaller companies, using an IDC service is the only way that they can have a successful presence on the Web. For larger companies, using an IDC (rather than an in-house data center) allows internal IT resources to be devoted to business applications and other projects more aligned with the core business. Regardless of the size of the user, selecting the right IDC is critical for success, and this paper has offered suggestions on some of the important criteria to use.

The development of the IDC market in Japan was profiled, and market data and typical prices presented. Case studies of several service providers offering IDC services in Japan have been used to highlight the types of offerings available, the differences in the way in which service providers support users, and how the different offerings are presented in the marketplace. The case studies covered two large carriers: one U.S.–based multinational service provider and one smaller local firm.

The four case studies, although all for IDCs located in Japan, represent offerings typical of those in the industry in many countries (with perhaps pricing being the greatest variable country-to-country). Along with the more general tips, they offer useful insight to any enterprise or other organization seeking IDC services.

DSL across the Globe

Patrick Kiernan
Director, Market Management
Spirent Communications

While digital subscriber line (DSL) issues in the United States have been discussed at length, the status of DSL in the rest of the world remains less clear. Yet it is important to be aware of global DSL issues for several reasons. Vendors may want to consider the opportunities that might arise in international expansion. Service providers will find the environment elsewhere amazingly similar and can learn lessons from foreign providers' experiences. Moreover, service providers are expanding into Europe, South America, Asia, and Australia.

What lessons does the DSL experience elsewhere provide for these vendors and service providers? After some brief observations about DSL deployment worldwide, this paper will briefly cover the status of deployments in Asia and Europe to draw some useful conclusions about the issues that have an impact on successful DSL rollout. It will then offer some ideas for solutions to the problems.

DSL Worldwide

It might be surprising to know that the United States does not have the largest DSL subscriber base in the world, and it might be even more surprising to learn that Korea holds that distinction. Korea has more subscribers than the United States, and because the actual growth rates for the two are comparable, this difference is likely to continue. Overall, Asia is a larger market than North America. Together, North America and Asia have accounted for most of the volume of DSL growth worldwide. What Europe and Latin America have lacked in volume up to this point, however, will be made up for by the fact that the compound growth rate over the next five years in those two areas of the world will be much larger than the rates in Asia and the United States. This situation is clearly the result of early entry into the DSL market in Asia and North America.

Deployments in Asia

The earliest deployment of DSL was in Singapore in 1997, so the history of DSL deployment begins in Asia (see *Figure 1*). That first deployment was followed by deployments in Taiwan and Hong Kong, both in 1998. Both of those deployments offered very broad DSL coverage. Korea entered the market next, and the service really took off from there. Indeed, although it was once the largest deployment outside the United States, as *Figure 1* indicates, it has now taken over the lead.

Thus Asia is home to the most successful DSL deployment to date. In both Korea and Taiwan actual DSL deployment has gone far beyond even the aggressive industry forecasts, and it continues to accelerate. The key driver in the Asian market has been high-speed Internet connections for gaming and video on demand. Yet rollouts in Japan and China remain important unknowns.

Japan is something of a wild card in DSL deployment. It has a high standard of living, and its people are very interested in the latest technologies. So it was expected that the Japanese would be very attracted to broadband offerings, and a big push was anticipated in 2001. It has not happened yet, however, because of a wide range of problems, including pricing, regulatory, competitive, and business-management issues. The Japanese national telecommunications company has made a huge investment in fiber to the home, a competitive technology, and pricing, infrastructure, and regulation make it difficult to bring in a second carrier to offer DSL.

The second-carrier approach, by the way, has proven successful elsewhere. In Australia, the incumbent carrier had not offered DSL aggressively. Because the telecommunications market there has been heavily deregulated, five or six competitive carriers were able to meet market demand for DSL successfully.

The Chinese market is another Asian wild card. The potential market is huge, offering almost boundless opportunities. The issues here are not related specifically to the telecom industry; instead, issues relating to government regulation of culture are making expansion difficult. Fundamental questions about whether a billion people will have free access to information are at the basis of delays in deploying DSL. The answers to those questions will determine whether DSL will be a skyrocketing success story in China or whether another era of repression will prevent it from taking off. Another factor in the situation is the decision to hold the 2008 Olympics in Beijing. The Olympics will increase opportunities for freedom in China and prompt development of the telecom infrastructure.

Thus China and Japan offer major potential markets for broadband services, but whether the potential will be realized remains an unresolved issue.

FIGURE 1

Deployments in Asia

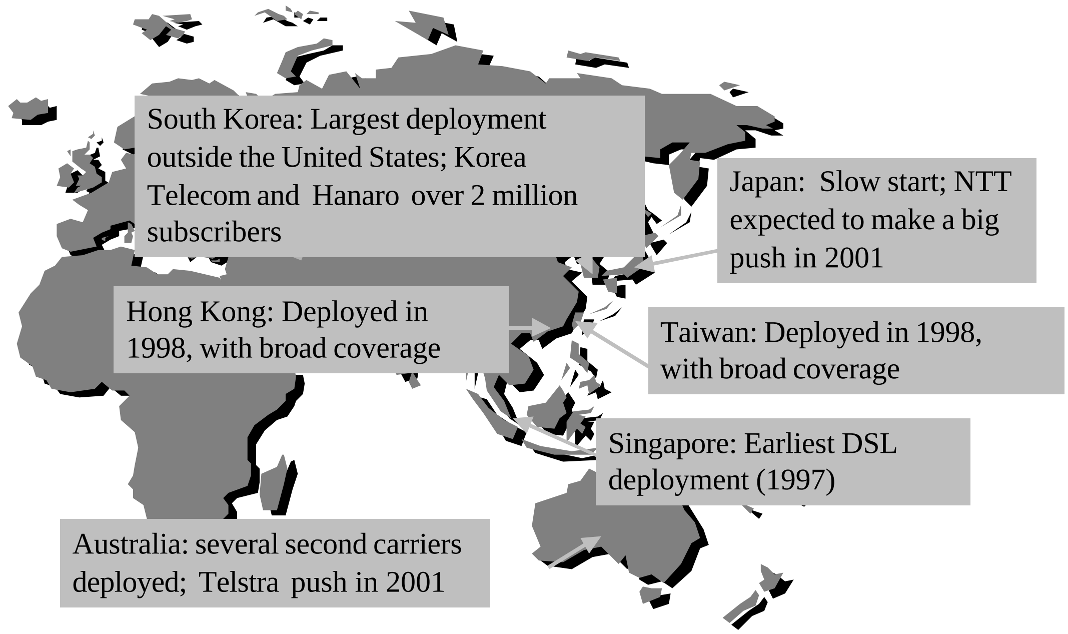

Deployments in Europe

Europe was somewhat slower in rolling out DSL—all the major incumbent local exchange carriers (ILEC) began to offer it in 1999 or 2000, slightly behind the rest of the global market. As in Australia, Europe has a significant amount of second-carrier activity. Regulatory issues throughout Europe, however, make it unclear whether those carriers will survive.

Generally speaking, the addressable market in Europe is larger as a percentage of households passed than in the United States. The European providers benefit from two key market characteristics: good coverage by carriers and very short loops. Central offices (CO) tend to be very small because something between a remote terminal and a full CO, called a controlled-environment vault, is deployed in many places. Use of this type of facility enables many carriers to enter the market with a small footprint. But they still have to be able to do a good job of providing coverage.

Another European characteristic is the impact of the European Union's requirement for local-loop unbundling. This requirement should introduce more competition and drive demand. Finally, scaling networks remains an issue in Europe.

Issues Affecting Successful DSL Rollout

Regulatory Issues

Despite European Union doctrines requiring local-loop unbundling, competitive local exchange carriers (CLEC) in Europe are complaining about the delay tactics ILECs are using to slow access to CO infrastructure. In the United Kingdom, for example, British Telecom is not complying, and competitors are pulling out of the market there as a result. In other instances, although the ILEC may be complying with local unbundling regulations, in fact the market is not open. In Germany, competitive carriers cannot afford to build a footprint or put equipment in a CO. In France, competitive carriers cannot enter only the COs in the area they want to serve; they must have a presence in all 3,000 COs, making it price prohibitive for a competitive carrier to establish a footprint. Despite regulations, then, it is not always practical to roll out DSL service.

Loop and Data Network Problems

DSL maintenance testing provides some interesting challenges, especially in Europe (see *Figure 2*). About 25 percent of the European physical plant is built on aluminum, not copper, which makes quality an issue. Another issue has emerged in the United Kingdom. British Telecom (BT) just lowered the price for wholesale DSL, thus eliminating the business case for competitive carriers to deploy a network. Competitive carriers become totally reliant on BT for every aspect of DSL, including testing. If BT does not test well, the competitive carrier does not know what is happening in its market. Many CLECs must purchase their own test systems, adding a level of expense, because they cannot trust BT testing. Other problems emerge from relying on the ILEC network. Not all problems reside on the loop; many of them reside on the data network, and the CLEC must have the tools to localize faults that exist in the data network. If the CLEC does not have access, it will not have that capability. Moreover, the ILEC is always in a better cost position than the CLEC.

Scaling

Scaling raises other issues. In Europe current operational procedures are completed manually, and telecommunications

FIGURE 2

Loop- and Data-Network Problems

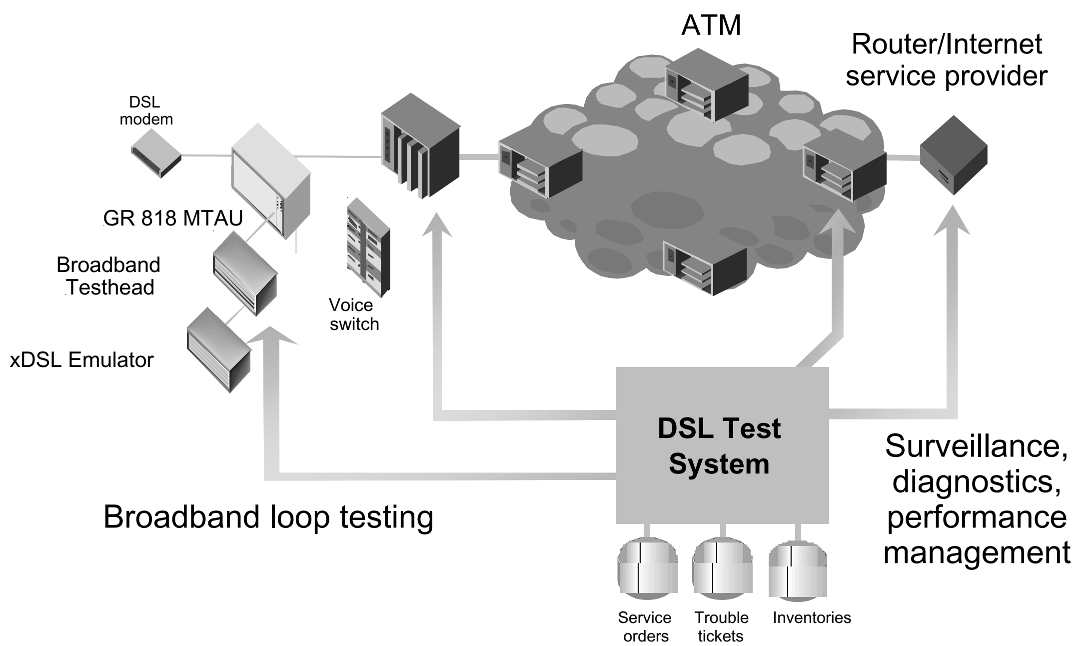

unions for the workers responsible for those procedures are backed by the government. It is difficult to change the manual processes, which are resulting in increasing backlog and thereby placing a cap on growth. An associated problem is a shortage of high-quality, well-trained technicians, making it even more difficult to carry out the required manual processes.

If growth is to occur, costs must be managed. Truck rolls must be reduced, the business must grow faster than the head count, and time to revenue must be shortened.

Solutions

The solutions to the issues that have affected DSL rollout abroad can be summed up in four points: (1) good pre-qualification tools to expand the addressable market; (2) self-installation; (3) remote testing; and (4) flow-through operations support systems. These lessons, learned from international DSL implementations, are important for DSL in the United States as well.

The Strategic and Financial Justifications for Convergence

Mike Kisch
Marketing Communications Manager, Enterprise Voice and Video Business Unit
Cisco Systems

Bill Erdman
Director of Marketing, Enterprise Voice and Video Business Unit
Cisco Systems

The intent of this white paper is to assist senior managers in defining the information technology (IT) and networking directions of their companies and to provide the necessary strategic and financial justification to make effective decisions regarding voice, data, and video convergence.

The term *convergence*, also known as multiservice networking, refers to the integration of data, voice, and video solutions onto a single, Internet protocol (IP)–based network. The technologies and tools for companies to converge their networks exist now as evidenced by a recent META Group study that found that 26 percent of Global 2000 enterprises are already in the process of migrating to a converged network (see *Figure 1*).

Convergence has progressed to the point that most organizations should seriously evaluate its role in the future of their networks. A converged network can play a critical part in helping a company identify new ways to generate revenue, reduce operational costs, increase organizational flexibility, and generate a sustainable competitive advantage. Many of the new business applications that are now deployed on converged networks provide immediate ways to increase personal and workgroup productivity while enhancing customer care and responsiveness. Convergence can also accelerate business cycles and help enterprises to realize the benefits of their IT investments faster.

As an IT investment, convergence is unique because of its ability to impact the entire organization. Whether it is workforce optimization, e-commerce, supply-chain management, or customer-relationship management (CRM) initiatives, a converged network provides the necessary foundation to decrease implementation time and maximize an organization's investment in new technologies. This is especially important given that the costs to implement a packaged software solution is usually two to three times the cost of the actual software. Convergence also allows organizations to challenge many of the traditional assumptions regarding labor, facilities, and capital investment by providing the tools to empower the organization to mobilize its workforce and centralize its infrastructure and network management. However, the far-reaching implications of convergence often make it a difficult concept for organizations to understand and evaluate.

Moving to a converged network can substantially reduce an organization's total cost of ownership (TCO) for its network and reduce the ongoing costs required to maintain and upgrade the network. Also, by simplifying the administration of the network, an organization's IT staff can focus more on strategic initiatives that can generate demonstrable business benefits. While the hard cost savings are enough to justify migrating to a converged network, the less easily quantifiable "soft" benefits provide the most compelling argument for why an organization should begin migrating to a converged network.

For organizations to understand the reasons to migrate to converged networks, they must understand the impact upon not only a company's cost structure, but also how convergence can help their businesses more quickly and more effectively reach their overall goals.

Situational Analysis

IT projects are now faced with a greater degree of scrutiny than ever before. Recent economic fluctuations have made it that much more important for IT managers to justify the strategic and financial value of proposed IT investments. However, regardless of the economic environment, organizations must continue to invest in IT projects that support specific business goals and can produce quick returns. A well-planned IT strategy will allow a business to emerge from adverse market conditions in a better position to take advantage of new opportunities.

As IT has gained greater acceptance as a strategic component of running a business, the parties involved in making the decision have expanded to include business and financial managers who want greater strategic and financial justification of why they should invest the organization's capital. Business and financial managers want to understand how these investments will help them to realize their business goals of generating more revenue, reducing operational costs, and creating a sustainable competitive advantage in

FIGURE 1

Rate of Convergence—Global 2000 Enterprises

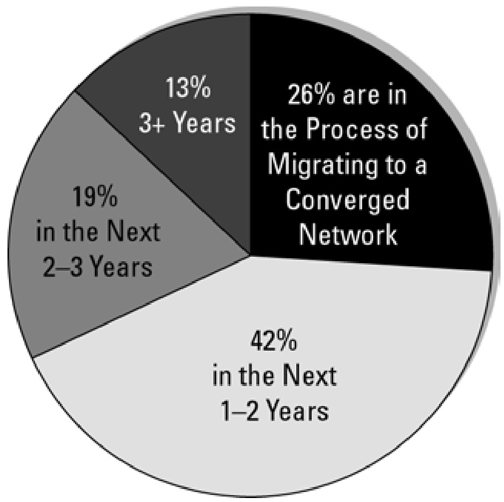

the hopes of either boosting or creating value for shareholders. They also want to understand how an investment in IT can accelerate the pace of an organization and increase organizational flexibility. According to Forrester Research, 68 percent of global enterprises believe that their network is a source of competitive advantage for their business.

Much has been written about convergence. As a result, it is often difficult to separate the hype from the reality. While most agree that converging to a single IP–based network is inevitable, there are divergent opinions as to when this will happen. A recent survey by the META Group of Global 2000 enterprises suggests that many companies are well on their way to evolving their networks. In fact, 73 percent of Global 2000 enterprises are already migrating to a converged network or are planning to do so within the next one to two years. This statistic supports the notion that the technology to converge disparate voice, data, and video networks has transitioned from the early adopter to early majority stage in the technology life cycle. Early adopters of IP telephony have proven that voice over an IP network is reliable, even across shared low-bandwidth circuits. Sending voice over the wide-area network (WAN) was one of the initial drivers behind migrating to a converged network. While recent decreases in public switched telephone network (PSTN) tariffs have closed the gap slightly, there still exist savings for large multinational companies that generate substantial costs for international calls. Initial deployments of IP telephony have ranged from several hundred phones all the way up to 25,000, highlighting the scalability of a converged network.

IT decision makers are faced with the task of creating a vision for the future of the network that will help their organizations to achieve their business objectives. At the same time, organizations must also identify the correct stage in the technology life cycle at which the best benefits can be realized while incurring the least risk. The technology life cycle plays a critical role in the investment decision and is often the most difficult aspect of evaluating an IT investment. Many organizations will not invest in a new technology until it has been proven and commoditized. This approach reduces the risk associated with the IT initiative, but even more significantly reduces the potential benefits. Assuming that an IT investment must be a key source of an organization's competitive advantage to be justified, this is only true if the technology is adopted and integrated faster and better than the competition.

The critical drivers that should lead an organization to begin evaluating converging their networks are the hard cost savings related to infrastructure, staffing, and facilities, as well as the improvements in productivity and customer care. There are, however, several key transitional events that can accelerate the evaluation and adoption process, as well as the financial payback:

1. Building a new office or moving to a new location
2. End of lease for private branch exchange (PBX) or support contract
3. Necessary upgrades for data network
4. Lack of expansion capacity of current voice network

In addition to these events, convergence makes particular sense for organizations that are rapidly changing in size and location by opening new branch or regional offices. Adding further urgency to the development of a convergence strategy is the exponential growth of data traffic relative to total network traffic. Data traffic is now growing 10 times faster than voice traffic and will surpass total voice traffic in 2001. As a result, most enterprises will need to place more of a focus on and make significant investments in new data infrastructure over the near term to meet network demands. As they evaluate new data equipment purchases, they should also ensure that the networking equipment could support not only data, but also voice and video.

Scope of the Analysis

Enough empirical evidence now exists to truly evaluate the financial and strategic implications of convergence. This analysis is intended to be an accurate comparison between the currently separate data, voice, and video networks and a converged network. Until now, the comparison most commonly made has been between a converged network carrying voice and the PBX. This is not an accurate comparison, given the greater functionality of a converged network, and is equivalent to comparing the personal computer (PC) to a typewriter. Both the PC and the typewriter are capable of word processing, but word processing is all the typewriter can do, while the PC increases productivity exponentially, enhances customer care, and distributes knowledge throughout the enterprise. An organization needs to evaluate the costs and benefits of their currently separate networks (voice, data, and video) with those of the converged network. Only then will the true benefits of convergence come to light. The majority of the benefits in this analysis focus on hard cost savings derived in the areas of equipment, staffing, and facilities. In addition, the analysis considers business benefits related to the deployment of new applications that can improve productivity and enhance customer care. When performing a cost/benefit analysis, it is very difficult to quantify these soft benefits. As a result, the return on investment (ROI) may significantly understate the true business value of convergence.

The Future of the Network

The future of the corporate network is impossible for anyone to predict. What can be reasonably expected is that the network will continue to become a more critical source of a company's competitive advantage. According to a recent survey, 92 percent of IT professionals believe that future IT initiatives need to support specific business goals. These goals include reducing the cost of doing business (58 percent), increasing business productivity (55 percent), and improving customer service (53 percent) (see *Figure 2*).

The ability of an organization to leverage IT investments that accelerate its business processes, take better care of its customers, and maximize the output of their physical and human capital will play a large role in determining who will continue to be successful. The demands that are placed on the network will also rapidly increase. According to Forrester Research, 77 percent of companies expect their Internet bandwidth usage to increase by more than 60 percent during the next two years. Such dramatic increases make it almost impossible to effectively plan capacity expansion with today's current network model. These increases in bandwidth usage will be driven by an increase in new users, a shift toward richer forms of communication that incorporate audio and video, and an even greater focus on e-commerce initiatives. Some of these e-commerce initiatives will be focused on increasing the current one percent browser-to-buyer conversion rate through the use of multimedia on-line collaboration tools that seamlessly integrate voice, video, and data and are designed to provide the customer a more satisfying on-line experience.

The IT investment process will also change as greater pressure is placed on IT decision makers to financially justify any new IT investment. An IT investment should no longer be approved unless it can help a company quickly generate new revenue streams, reduce operational costs, or create a sustainable competitive advantage (superior customer service or faster speed to market). This need for IT initiatives to rapidly contribute to the company's strategic business goals highlights a critical limitation of the existing network model. Separate voice, video, and data networks require separate sets of infrastructure, multiple teams of administrators, and constant concurrent maintenance. These networks have evolved separately, driven by isolated communication requirements for voice and data. As a result, these separate networks are extremely complex. According to Forrester Research, 90 percent of businesses see this complexity continuing to increase because of more difficult-to-manage components, more users, and more network links.

Many of the IP–based, converged applications discussed in this analysis can also be deployed over separate voice, video, and data networks. The primary difference is the ease and speed in which these applications can be deployed, maintained, and integrated into the network infrastructure. One of the primary benefits of IP is that it offers a transmission protocol in which data and voice communications can be combined together in an integrated fashion, reducing network complexity and providing an enriched user experience. This inherent complexity highlights one of the primary limitations of the existing multiple-network model. It is a difficult and costly process to integrate new applications and technologies. A good example is unified messaging (UM). UM has existed for more than five years and was originally developed for separate voice, video, and data networks. Initially, it showed great promise, but has since failed to achieve its user adoption projections. No one disputes that UM has the potential to increase employee productivity and reduce IT support and administration costs. Unfortunately, the cost and time necessary to integrate UM into disparate networks often overshadow the benefits. Another relevant example is computer telephony integration (CTI). CTI was developed to bridge the gap between the data and voice networks and initially showed a lot of promise. However, it proved to be so complex and expensive to implement that it could be financially justified only in a call-center setting. A converged network significantly reduces this cost of integration and promotes the deployment of applications at the department and individual levels.

Business Impact of Convergence

Convergence can lower an enterprise's total cost of network ownership through the elimination of multiple sets of infrastructure, simplified administration/maintenance, and consolidation of IT staffs. However, the real benefits are driven by the new services and applications that are more easily deployable on a converged network than on three separate networks. For the purposes of simplifying this analysis, these business benefits have been grouped into four benefit categories, Business Empowerment, Personal Productivity, Workgroup Productivity, and Customer Care.

Business Empowerment

Convergence can empower a business to reduce infrastructure, staffing, and facilities costs. It can also increase organizational speed (integration of new acquisitions or offices, and deployment of new applications) and flexibility (ability to quickly react to changes in the market).

FIGURE 2

Objectives of IT Investment

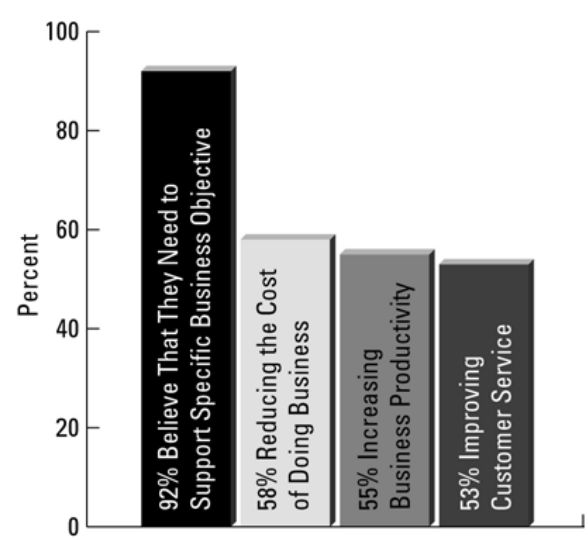

Reduction in Network Infrastructure Costs

As voice, video, and data networks continue to evolve to meet the pressures placed upon them, they grow significantly more complex and expensive because organizations must scale up three separate networks. A converged network can do the following:

- Reduce investment in multiple network infrastructures by converging to a single IP–based network. An organization will no longer need to invest in a dedicated device such as a PBX or maintain a separate integrated services digital network (ISDN) for its videoconferencing needs. The migration to single network will significantly reduce the complexity of an organization's network infrastructure, leading to easier administration, increased flexibility, and greater scalability.
- Lower system-expansion costs by deploying a single network, enabling scalability to meet new demands placed on the network by new users and new bandwidth-intensive applications. The resulting single point of administration and decreased reliance on outsourced services will also reduce ongoing network-maintenance costs.
- Centralize application hosting and call processing without sacrificing the user experience at the branch office. This allows remote offices to be offered the same applications as the main branch without having to invest in their own infrastructure and software. This will also provide the central office (CO) a greater degree of cost control over what is added to the network, ensuring better systems integration and security.
- Reduce the number of wiring drops per user by 33 to 50 percent in a new facility, which results in a fast payback for a new investment in converged network technology.
- Significantly reduce the costs associated with integrating new services and applications. For example, UM in a multiple network environment is burdened by the need to tie into proprietary devices, such as the PBX, that have their own unique protocols. This can lead to delays in implementation and therefore delays in ROI. In a converged network, UM can also help a company reduce the number of distinct directories that it needs to maintain by integrating voice, e-mail, and fax into a single lightweight directory aAccess protocol (LDAP) directory. This is another significant source of short-term savings, because it is estimated that some IT groups spend approximately 25 percent of their time per year maintaining and updating the many directories in their networks.
- Lower hardware-connection costs within an organization by adopting an IP–based network. For instance, to connect an existing voice-mail server to a PBX requires a channelized T1 line with 1.5 megabits of bandwidth that can support 24 users and costs approximately U.S.$6,000. In a converged IP–based network running 100 megabit Ethernet with a UM solution, a single server can support a similar number of sessions at a cost of only U.S.$600. In other words, for 1/10th the cost, an organization is able to increase its performance 100x.

Reduction in Staffing and Administration Costs

A recent survey by CIO determined that 49 percent of businesses believe that an IT skills crisis exists. As a result, companies are having to rapidly increase salaries to retain qualified IT staffers while increasing training budgets by more than 25 percent to teach them the necessary skills to manage the growing complexity of the corporate network. A converged network can reduce IT staffing costs and simplify network administration by doing the following:

- Consolidating IT staff skills sets. A converged network based on IP reduces the need for network specialists to manage specific technologies, such the PBX, thereby allowing an IT organization to hire a more generalized skill set with a common knowledge base and be better able to handle growth of the network or departures of key IT personnel.
- A converged network will allow an organization to perform more critical IT functions internally as opposed to having to outsource them providing them greater control over their network and faster response time to their users.
- A converged network will allow organizations the ability to manage much larger user communities with smaller staffs. New converged networks support centralized application hosting and call processing, thereby reducing the need to have dedicated IT staff at each individual work site.
- Moves, additions, and changes account for as much as 14 percent of an annual IT department's budget. Each move is estimated to cost an organization between U.S.$75 and U.S.$135. The average company moves its employees approximately once per year. The potential cost savings by allowing employees to move and install their own phones can be significant for companies of all sizes.

A converged networking model is also much more capable of supporting today's increasingly mobile workforce. This is especially important for salespeople, consultants, telecommuters, or computer technicians who spend much of their time traveling from location to location but are unable to have access to the same network capabilities that they have in their office. During the past several years, enterprise mobility has become a commodity as a result of the adoption of dynamic host control protocol (DHCP). IP phones can effectively leverage this technology, resulting in cost savings and increased employee productivity.

Reduction in Facilities Costs

A single converged network can also help a company to optimize its use of existing facilities by providing a more flexible foundation to serve the needs of its users.

- Most companies have employees who are constantly on the road and rarely use their offices. Unfortunately, other employees cannot use these unoccupied office spaces, because their phone extensions do not travel with them. The practice of hot-desking, or hoteling, has been in existence for almost a decade and originally promised significant savings and greater productivity from investments in real estate and facilities. However, many organizations found that the cost to adapt their current networks to accommodate mobile employees and the associated inconveniences outweighed the benefits. A converged network with IP phones enables employees to freely move about the

office to available slots, allowing companies to assign more workers to shared workspaces and maximize their ROI in facilities.
- Convergence also allows companies that are rapidly expanding or contracting because of mergers or divestitures to quickly integrate new sites into corporate networks.

Convergence can significantly reduce a company's TCO of their network. But to fully understand these cost savings, IT executives must modify their evaluation processes to give convergence its fair due. An organization contemplating migrating to a converged network must also understand how a single network can improve its ability to scale for the future and quickly react to the dynamically changing needs of the business. Specific types of change can include the deployment of new applications and services designed to improve business processes or enhance customer care, the addition of new users to the network, and the integration of new hardware solutions. A single converged network also simplifies the ongoing network-design and evaluation process, allowing an organization to speed up the implementation of new technologies. It is difficult to assess the monetary value of these last several intangible benefits, but all IT managers understand the importance of having a network that can empower change and innovation rather than impede it.

Enhance Personal Productivity

Perhaps the most significant benefit of a converged network is the array of applications that can be quickly integrated into the network and deployed. These applications can enable organizations to increase employee productivity by streamlining administrative tasks, allowing them to focus on activities that create new revenue streams or generate cost savings. As mentioned before, one of the primary drivers of convergence is its ability to provide a flexible foundation to easily integrate all types of new applications and services. Improving personal productivity can also be a key driver in reducing operational costs, because it can help an organization maximize the output of its existing labor force without the need to hire new people. This is equally important during downturns, when there are often hiring freezes, and during upturns, when the competition for employees can drive salaries sky high.

The power of the networked world is in its ability to allow people to more rapidly communicate through multiple delivery mechanisms. Voice-mail, e-mail, and fax have greatly enhanced the ability to communicate across an organization and with customers. However, it has also created a staggering amount of inefficiency in the enterprise. Employees now spend an average of 2.5 hours per day listening, reading, and responding to voice-mail, e-mail, and faxes (see *Figure 3*). Identifying ways to help employees to more efficiently manage these communication methods will be the source of significant productivity improvements.

Unified Messaging

A study done by The Radicati Group, Inc. found that UM can generate 25 to 40 minutes of additional productivity per employee per day and can reduce IT support and administrative costs by 70 percent. UM provides users the ability to access and immediately respond to voice, fax, and e-mail messages from any phone or PC within the enterprise, reducing the time associated with accessing multiple devices. In contrast to today's voice, fax, and data-messaging systems, in which content must be manually copied or scanned to be passed between different system types, UM on a converged network supports a universal inbox that can contain all three types of messages. Employees can access voice-mail, e-mail, and faxes while traveling, and then will be able to reply or forward to the appropriate person (see *Figure 4*). As mentioned earlier, UM has existed for more than five years and has demonstrated tangible productivity benefits. Yet it is still a relatively undiscovered technology because of the difficulty to integrate the software into a multiple-network environment. A converged network provides the necessary platform to make UM a reality for many organizations.

Personal Communications Assistants

Because employees typically have multiple contact points, it is becoming increasingly difficult to know the appropriate number to call at a given time. As a result, "phone tag" is rapidly increasing, leading to missed calls and a frequent inability to contact a critical resource in times of need. In addition, it is often necessary for workers to prioritize who can contact them at a given time. Personal communications assistants on a converged network resolve these problems by providing employees tools to forward critical calls to their most appropriate number, to screen and prioritize incoming calls, and to easily set up audioconferencing bridges without the assistance of an outside party (see *Figure 5*). Some personal communications assistants also have speech-recognition capabilities that allow users to dial by name. Today's personal communication assistants also allow employees to customize call-screening and forwarding features through a graphical user interface (GUI) without having to request assistance from someone in the IT department.

FIGURE 3

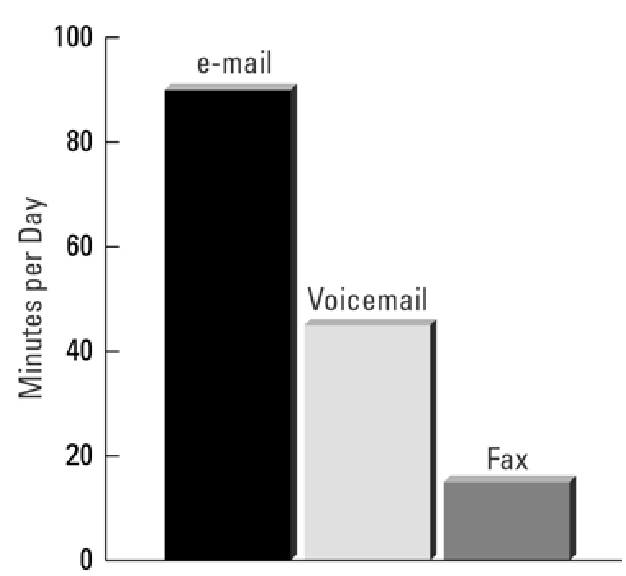

FIGURE 4

Unified Messaging

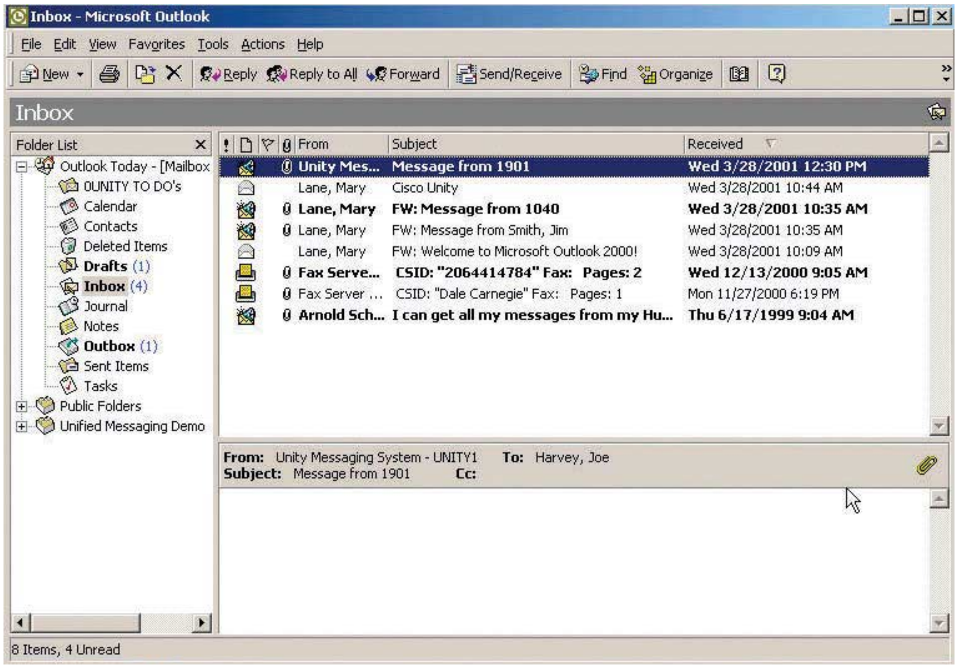

IP Video Solutions

Videoconferencing has also been in existence for many years but has never gained widespread adoption, largely because of the cost of having to acquire and maintain a completely separate network to handle video. As a result, when companies did invest in the technology, it was often only for use by a select group of individuals. A converged network has the ability to finally put the power of videoconferencing in the hands of everybody by providing companies a more cost-effective and more easily deployable model. The primary benefits of videoconferencing are its ability to save on travel costs, to minimize lost employee productivity due to travel, and to provide a richer form of communication among people at different locations. Similar to UM, it has been held back by the limitations of the traditional network model. In a converged IP network with QoS, an organization can provide videoconferencing and video on demand (VOD) capabilities to the desktop. New uses for this technology could include distance learning, whereby employees can access video content at their own leisure, as opposed to having to travel to a central training facility, and easier dissemination of critical business communications, such as quarterly board-meeting updates. Finally, IP video solutions have the power to further enable the mobile workforce by providing remote employees a real-time face-to-face interaction with their office-bound colleagues.

IP Phones and IP Soft Phones

IP–based phones are intelligent communications devices capable of supporting many new features that can also increase personal productivity. Many of these new IP phones use extensible markup language (XML)–based applications to allow the user to view employee directories, daily calendars, e-mail messages, and voice-mail using pixel-based liquid crystal displays (LCDs). The advent of the XML programming language allows organizations to easily download critical information and customized applications to a user's phone. An IP phone can provide users a potentially faster alternative to the computer to access simple bits of information that will speed up their work processes. Opportunities for this type of technology will develop on an industry-by-industry basis as organizations identify ways to deliver customized information to their employees or customers.

New PC communications applications currently allow users to eliminate their hard phone and use their PC for voice calls. A simple-to-use interface eliminates the need for two devices and allows the user to easily set up conference calls with the touch of a mouse. In addition, because an IP soft phone can travel with the user, telecommuters with high-speed access will be able to take their phone extension home with them and receive calls just as they would in the office. A converged network and the applications enabled by it can significantly improve the performance of employees by providing them the tools that they need to work smarter and faster.

Improve Workgroup Productivity

Even in today's information society, most knowledge still resides in the heads of a company's employees. Companies looking to the future are trying to identify methods to tap this individual specific knowledge and share it with the entire enterprise. As the workforce becomes more mobile, this problem will continue to increase. A converged network offers a platform that allows the deployment of applications and services that are intended to allow an easier dissemination of this knowledge across the organization. Several existing applications can play a key role in this endeavor, as are described in the following sections.

FIGURE 5

Personal Communications Assistant

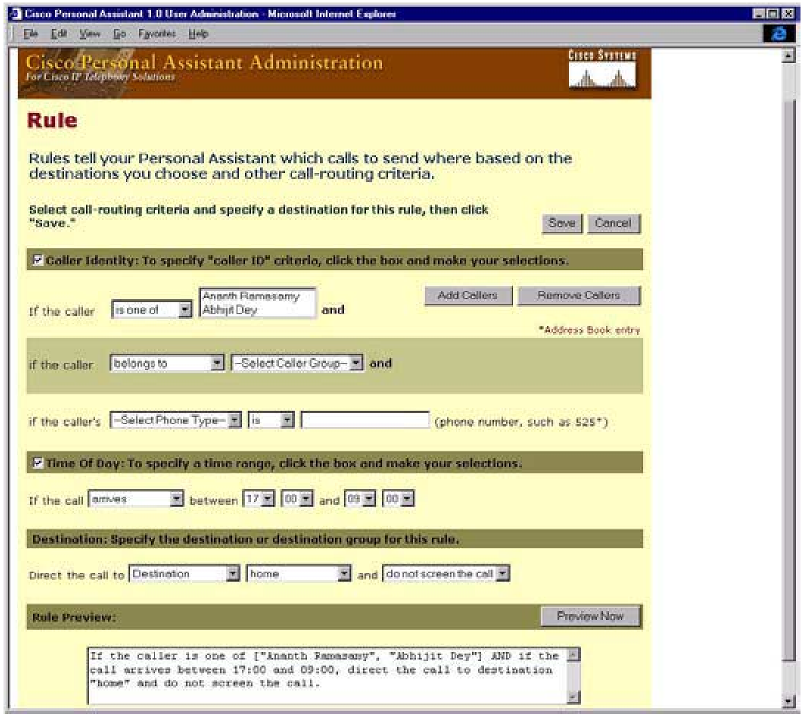

IP Video Solutions
Videoconferencing has always had tremendous potential to improve knowledge sharing across an organization. But, as mentioned before, only a small percentage of people within an organization were allowed to use it because it was expensive and usually located in only one conference room. For videoconferencing to be a viable method to distribute knowledge, it has to be available to all employees. Under the convergence model, videoconferencing will be transmitted as IP packets and will be able to utilize the same network as does voice and data traffic. This will allow a much wider usage of the technology in the organization. Multiple groups of employees separated by thousands of miles will have the ability to collaborate on projects and share their knowledge, thereby increasing the productivity of the entire team, shortening the decision-making process, and hopefully accelerating the organization's time to market. VOD also offers the ability to provide geographically dispersed users access to an expert's knowledge through distance-learning programs.

Collaboration Tools
Collaboration applications that combine voice and video interaction with information sharing will play a key roll in improving knowledge management by facilitating easier access and the exchange of critical corporate information. This ability to share knowledge across an organization will allow a business to more quickly recognize problems and generate solutions or competitive responses. Current features include application sharing, whereby two users are able to read through a document together to agree on the correct content. Additional applications of the technology could include on-line editing of video content, walking a remote group of people through a presentation, or sharing a white board to brainstorm ideas while in multiple locations. The convergence of voice, video, and data onto a single network helps to make the obvious benefits of collaboration tools a reality.

Enhanced Customer Care/Responsiveness
As evidenced by the boom in CRM, it is obvious that most companies finally understand the value of keeping their current customers happy. An oft-cited statistic that continues to resonate is that it is 5x to 10x more expensive to acquire a new customer than it is to retain an existing one. Many of the recently deceased business to consumer (B2C) companies learned that lesson the hard way as they poured outrageous amounts of money into campaigns that were intended to turn browsers into buyers but often failed in their ability to build a loyal customer base. Most people also recognize that superior customer care and responsiveness is a primary driver of competitive advantage. Every industry has an example of one company that excels based upon its ability to satisfy the customer better than its competition.

Superior customer care has been a particularly vexing issue for companies who are in the process of major e-commerce initiatives. How can a company create an on-line environment that can provide a comparable level of service to a physical store staffed by real people? Converged networks provide organizations a fully integrated voice, video, and data infrastructure designed to provide their customers a richer on-line experience.

Multimedia Contact Centers
Converged networks can provide an enterprise the opportunity to create a truly multimedia call center. In the past,

call agents have only been able to interact with customers via the phone. Customers would at times be frustrated because they were unable to see the person to which they were talking and, oftentimes, they were forced to provide simple account data multiple times before their problem could be resolved. A multimedia contact center will allow an organization to put a face with the voice, and it will allow easier access to data critical to dealing with the customer's problem as fast as possible. The converging of data, voice, and video onto a single network will also provide the Web browser the option of click-to-dial functionality, where they can click a button while on the Web site and be put in contact with a customer-service agent. Furthermore, the easy access to customer data can lead to cross-sell and up-sell opportunities for an organization.

The enhanced mobility features of a converged network will allow call agents to be located remotely. This can reduce labor cost and increase employee quality by allowing the organization to recruit on a regional or national basis rather than within one specific market. It can also reduce the investment in call-center facilities, and the flexibility provided to employees could potentially boost employee morale.

Collaboration Tools

Collaboration applications and their ability to directly affect the customer's user experience are one of the most compelling reasons to justify migration to a converged network. With more than two-thirds of all on-line shopping carts never making it through checkout, there is an obvious disconnect between the level of service that customers receive and the level of service that they require to be successfully converted from a browser into a buyer. This disconnect cost North American businesses more than $1.6 billion in e-commerce revenue last year. And, according to Jupiter Research, more than 90 percent of on-line shoppers want some form of on-line human interaction. Collaboration software running over a converged network will allow the customer-service agent to look at the same page on the Web site as the customer, to make recommendations and product comparisons, to instantly transfer relevant documents to the user's desktop, and to walk them through the checkout process. As customers continue to demand a level of service similar to that which they receive in a physical store, tools such as on-line collaboration can provide the customer a customized on-line experience that eliminates many of the dehumanizing qualities of using the Internet to shop or seek help for a problem.

Contact-center applications and on-line collaborative tools offer the most obvious customer-care benefits. However, UM and personal communications assistants can also improve customer service by providing customers quicker response times to their inquiries. For instance, if a customer urgently needs to reach their account manager, personal communications assistant software allows that call to be routed to the salesperson's current location rather than his or her voice-mail box. Similarly, UM reduces the time required to access and read or listen to voice-mail, e-mail, and faxes. This, in turn, can reduce the time that it takes to respond to a customer.

Migrating to a converged network and the applications enabled by it can not only offer an organization short-term reductions in its TCO, but also can generate substantial long-term benefits, such as increased personal and workgroup productivity, enhanced customer care, and increased organizational flexibility. In other words, it has the potential to be a source of strategic value rather than just a cost center.

Issues That Arise When Migrating to a Converged Network

This analysis has explored the benefits of migrating to a converged network. However, there are also potential issues that could slow adoption of this technology. According to a recent study by the META Group, the primary obstacles impeding organizations from migrating to converged networks are perceptions about voice quality, system reliability, interoperability with legacy systems, and cost.

Latency, jitter, and echo were early problems that plagued IP telephony. These were largely caused by a lack of QoS in the network. A converged network must be able to separate each traffic type and handle it according to its unique requirements. For example, data traffic is not time-sensitive—it travels in bursts and requires accurate delivery. However, voice—and to a certain extent, video traffic—is very time-sensitive. Adding voice packets to a bursty IP environment requires QoS to the desktop. An organization needs to understand what QoS is and how important it can be to ensure proper network performance for voice, in addition to data and video. Traffic classification and marking, queuing, and data-packet fragmentation and interleaving techniques are available now to guarantee voice quality. Planning a QoS strategy before deployment saves time and money and eliminates frustrated users. Most IP telephony vendors can now deliver toll-quality voice. However, to ensure this high quality, networks need end-to end QoS policy management in all routers and switches. Even if an organization is not fully committed to a convergence strategy now, it makes sense that all of its new data equipment can handle voice and video, because this will ease the path to eventual migration.

Reliability is also a critical concern of companies that are contemplating the convergence of their networks. Universal experience with traditional analog or digital phones is that, upon picking up the receiver, the user receives a dial tone 99.999 percent of the time. It is often assumed that when merging voice onto a data network, it will become unreliable. However, the PBX is inherently no more reliable than a data network; what makes it more reliable is that organizations recognize that voice is mission critical and therefore usually invest in the necessary redundancy and power back-up systems. Many IP telephony vendors have built similar levels of redundancy into their systems via call-processing server clusters, redundant routers and switches, and uninterruptible power supply (UPS) systems. With the correct design considerations and best practices, converged networks running IP telephony can achieve a comparable level of reliability to the traditional voice network. The added benefit of building redundancy into a converged network is that organizations can also improve the reliability of their data and video traffic. Another early problem with IP telephony was that the phones ceased to function when the network lost power. In response, most IP telephony vendors have developed in-line power solutions to ensure that users still have voice service during a power loss.

Most organizations have made significant investments in their existing voice, video, and data networks. They are understandably concerned about their ability to protect their investments while migrating these separate networks to a converged networking model. Therefore, a low-risk migration path is required from the old world to the new world. Most converged networking vendors have created products to ease this transition and ensure that new equipment can integrate with the existing infrastructure. In the longer term, a converged network will most likely make additional technology purchases more interoperable, because it will be based on open standards and will be one network as opposed to separate voice, video, and data networks. As in any technology investment, especially one with such far-reaching implications as convergence, an initial investment in the technology will be required to migrate to a converged network. For some organizations that are opening new offices, retiring PBXs, or already planning to make significant investments in data-networking equipment, the financial decision is relatively easy. For companies that have older data-networking equipment, the up-front cost of migrating to a converged network will likely be steeper. For these organizations, it is recommended that they begin to prepare their data network to handle voice and video before they make a substantial investment in IP telephony equipment and applications.

Summary and Conclusions

Convincing an executive committee to invest in new technology in these lean economic times can be a daunting challenge. The IT initiative must support core business strategies, yet must also provide quick returns to meet an organization's short-term financial focus. Migrating to a converged network can deliver both. Its ability to empower an organization to reduce infrastructure, staffing, and facilities costs can produce a quick ROI, while its greater flexibility and scalability allows a company to more quickly react to change. Finally, the ability to more easily integrate new end-user applications allows a business to improve personal and workgroup productivity and enhance customer care. All of these benefits can directly contribute to a business' competitive advantage, whether that advantage is based on speed to market, superior customer service, or being the low-cost producer.

To prepare for migration to a converged network, companies must understand all of the relevant factors to ensure a high probability of success. It is important to understand your existing PBX vendor's future architectural plans for your current voice network and how that vendor plans to provide viable migration paths to IP. Convergence can potentially take many forms, but some strategies are further developed than others. Also, to make migration to convergence easier, all organizations need to assess their current data networks to determine how well prepared they are to run voice and video over the same lines.

Lastly, IT managers must begin to evaluate convergence on its own terms rather than as merely a replacement for the traditional PBX. This will be a gradual process for some, but companies that immediately grasp the benefits of convergence will create networks that effect change and innovation rather than impede them. And if there is any doubt about the future of converged voice, video, and data networks, just look around the office and count how many typewriters you see.

The Consumer Relationship Is King: Thoughts on the 21st Century Content and Media Company

Michael Kline
Senior Manager
Deloitte Consulting

Overview: The Consumer Relationship Is King

"It's good to be the king!" (Louis XVI [Mel Brooks], History of the World Part I*)*

As content and media companies contemplate the strategies they need to forge their digital futures, one thing is clear: The rules of competition in the content and media industries are going to change. The impact of the transition from analog to digital content creation, aggregation, and distribution will be significant enough to require a new competitive framework to describe the ways companies should respond to this transition.

The purpose of this paper is to present our thinking on what that competitive/strategic framework looks like and what the implications are for content and media companies today. We intend this framework to address the concerns of the various players in the content and media value chain, from content creators to content aggregators to content distributors.

Here are some of the highlights of the framework.

- In the analog world, competition occurred along two dimensions: content and distribution. Content companies bulked up to ensure access to favorable distribution, and distribution companies bulked up to ensure favorable access to content.
- In the digital world, competition will occur along a third dimension: personalization. The battle for scale and advantage will increasingly occur along this axis.
- Personalization is a measure of the depth of knowledge a company has about the preferences of its customers and that company's ability to deliver customized products and services to its customers, based on that knowledge.
- For a number of reasons, we believe that today's personalization efforts will succeed, whereas those in the past—"mass customization", "one-to-one", et al.—were either prematurely introduced or inadequately implemented.
- Businesses historically powered by content and distribution assets will, over time, be increasingly powered by personalization assets, and value will flow disproportionately to those asset holders.
- Aggregation of consumers, as well as the information associated with those consumers, will create equal if not greater leverage in the supply chain than the aggregation of content.
- A new breed of company—the entertainment services provider—will emerge.
- Privacy will be a major issue for companies competing along the personalization axis.
- In light of these changes, content and media companies will need to develop competencies in personalization that cut across all aspects of their organizations.
- In particular, content creators need to begin moving from selling products to assembling packages of content and content-related interactive applications that enable them to capture end-user information and to pre-empt content distribution companies from capturing the value of complementary products.
- In the short to medium term, content owners will be king, as many content and media companies try to establish direct consumer relationships by paying large sums for content. In the longer term, consumer-relationship managers will be king, and the most successful content and media companies will have a superior understanding of the consumer and the ability to

deliver the consumer what she or he wants, to where she or he wants it, and when.

- While personalization will play a major role in the future, it is unclear who the leading consumer-relationship managers will be, and there are several different possible scenarios to describe the future evolution of content and media businesses. What is clear is that strategic, technological, and organizational flexibility will be required in order to respond to the various possible futures and to communicate to investors about the strategies employed amid this uncertainty.

The first section of the paper, "The Content and Media Competitive Landscape: The Analog View," describes the historical competitive framework that has characterized competitive behavior in these industries over the past decade. The second section, "The Content and Media Competitive Landscape: From Analog to Digital," examines the forces that are changing that competitive framework and that, in our opinion, require content and media companies to adopt a new framework to compete effectively going forward. The third section, "Personalization: The Third Dimension of Competition," outlines our thoughts on the new dimension of competition for content and media companies: personalization. The fourth section, "Strategic Uncertainty: The Possible Futures for Content and Media Companies," presents three different possible scenarios for the future evolution of the content and media industries and suggests the various strategies different companies will use to compete moving forward. The last section, "Owning the Consumer Relationship: Challenges for the 21st Century Content and Media Company," considers the organizational, technological, and strategic implications personalization will have for content and media companies.

The Content and Media Competitive Landscape: The Analog View

"We're going to need a bigger boat." (Chief Martin Bridy [Roy Scheider], Jaws*)*

The delivery of content to consumers has historically occurred along a fairly monolithic value chain made up of license holders, content aggregators, and consumer access/content distribution companies (see *Figure 1*). Competition between different segments of the value chain was less about fighting over end-user ownership and more about establishing leverage in negotiations between those segments. And, for the most part, each player in the value chain tended to focus on different customer groups, enabled by different assets and competencies, and powered by different economic models and motivations.

Content companies acted as venture capitalists, financing the development of original content (and the acquisition of third-party content). Over time, these companies diversified their risk by branching out into many different forms of content (e.g., film, music, publishing, games), all or most of which they sold/syndicated to content aggregators and distributors. Aggregators (e.g., television networks, cable channels, Internet portals) packaged and scheduled the content in order to maximize advertising revenue and, over time, branched out into financing original content creation in order to secure an ongoing supply of content and to establish their signature brands. Content distributors (e.g., cable and satellite companies, retail stores) leveraged heavy brick and mortar infrastructure investments to deliver an ever-increasing supply of channels/titles to end users via direct sales and/or subscriptions.

FIGURE 1

Content and Media —The Analog Chain

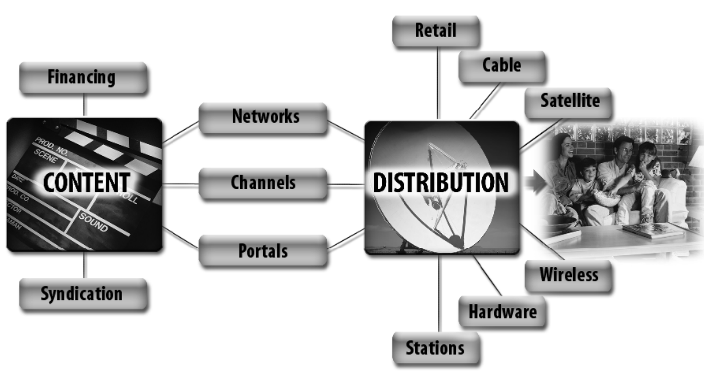

	License Holders	Aggregators	Customer Access
Primary Assets	• Licensed to breadth of content • Access capital	• Audience • Advertiser relationships	• Physical distribution assets • Direct customer relationship
Business Model	• Revenue share with distributors • Limited direct sales (clubs, etc.)	• Advertising • Limited subscriptions	• Subscriptions • Direct sales • Advertising
Strategy	• Aggregate breadth of content • Secure access to distribution	• Develop original programming • Aggregate content to secure distribution	• Consolidate distribution • Secure access to content

SOURCE:DELOITTE CONSULTING

In this analog world, with a single path to the end consumer, competition took the form of maintaining parity among the various value-chain participants. By way of analogy, content creators and aggregators, much like the United States and Russia during the Cold War, "bulked up" by acquiring other content creation and aggregation companies to ensure favorable access to consumers and advertisers. At the same time, distribution companies also consolidated in order to ensure favorable access to content (see *Figure 2*). In other words, competition tended to occur along two primary dimensions: content and distribution.

More recently, this wave of horizontal integration has been accompanied by a wave of vertical integration, where companies have made acquisitions to position themselves further along both the content and distribution competitive axes (again, see *Figure 2*). From Disney's acquisition of ABC, Viacom's acquisition of Blockbuster (and later, CBS), and the merger of AOL and Time Warner, the level of vertical integration in content and media companies increased dramatically during the last decade.

Firms did make different bets in terms of the level of vertical integration in which they engaged—e.g., Disney, for the most part, remained out of the "content distribution/consumer access" business in terms of cable and satellite system ownership (though it does have direct relationships via its retail and theme park operations), while AOL Time Warner integrated across the value chain. Likewise, certain firms made acquisitions favoring one revenue source over another—e.g., Viacom (advertising) versus Vivendi (consumer expenditures).

That being said, competition in content and media over the past decade has been, more than anything else, a battle for physical (and global) scale to maintain relative bargaining power in the value chain. Like the Cold War, however, the best outcome in this two-dimensional competitive framework is stalemate. Innovation often takes a back seat, which paves the way for new companies to enter. There is also the question of how much bigger content and media companies can become—partially, for regulatory reasons, but even more because of the diminished operational efficiencies and competitive nimbleness that come with size. But, most important, the advent of digital technologies poses a fundamental threat to the content and media value chain itself.

As we will discuss in the next section, the changes to the value chain that are being brought about by digital distribution technologies will, in our opinion, force content and media companies to re-evaluate how they compete.

The Content and Media Competitive Landscape: From Analog to Digital

"Size matters not. Look at me. Judge me by my size, do you?" (Yoda [Frank Oz], Star Wars: Episode V—The Empire Strikes Back*)*

The traditional (analog) content and media value chain was characterized by a single path to the end user (or, in the case of filmed content, single, well-defined schedules—or windows—for delivering content to the end user) defined by the largest players in the value chain. "However," as Thomas Middelhoff, the CEO of Bertelsmann, said, "it is my conviction that as content becomes digital, the entire media and entertainment value chain will change."

The digitization of content disrupts the traditional value chain by opening up multiple paths to the consumer of that content—that is, because digital content is no longer tied to the channel for which it is created. In the analog world, for example, music is essentially defined by the channel in which it is sold—i.e., indivisible units called compact discs (CDs) sold in physical retail stores—and, as such, music has only a limited number of ways of reaching consumers. In the digital world, music can be delivered to consumers through a variety of channels (e.g., peer-to-peer, virtual stores, personal

FIGURE 2

The Content and Media Competitive Landscape—The Analog Battle

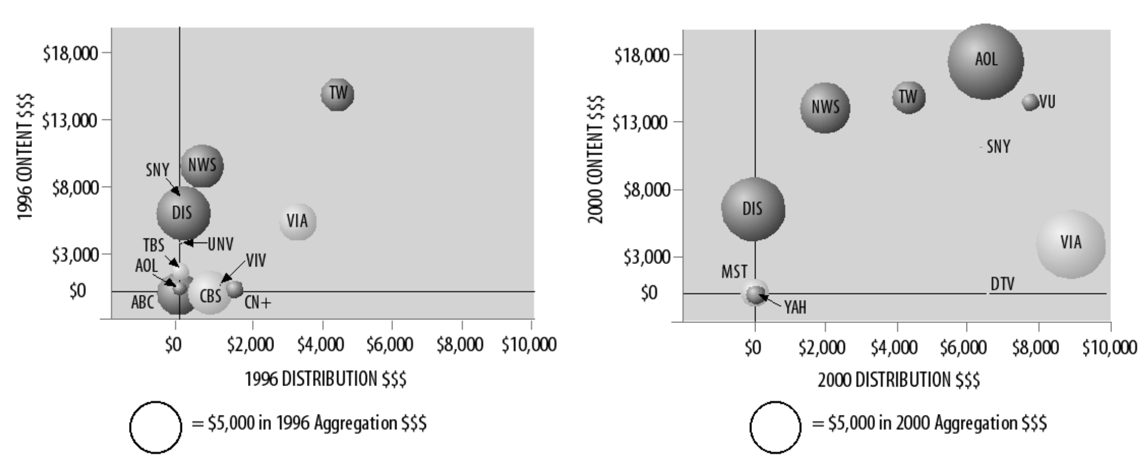

SOURCE: COMPANY ANNUAL REPORTS, DELOITTE CONSULTING ANALYSIS video recorders [PVRs]) and under a variety of different economic models (e.g., pay per download, monthly subscriptions—perhaps with incentives for contributed content). And while the existing music companies have successfully slowed that transition through legal action and scaled-down digital offerings of their own, the sale and consumption of music will change due to these technologies.

With digital technologies, the linear, analog value chain becomes a networked, digital ecosystem, characterized by many more points of contact with the consumer by many more players. In effect, what digital technologies are facilitating is the convergence of formerly unrelated industries and companies, with the locus of convergence taking place at the individual consumer (and advertiser) level. For example, the $22 billion home-video market is captured today by a limited number of players (e.g., video stores, video distributors, and the license holders). In a world of digital video on demand (VOD) (see *Figure 3*), many companies not historically involved in that revenue stream are positioning themselves to capture part of it. Companies such as TiVo and RealNetworks are attempting to become the networks of the future, and TiVo—at least for now—is able to capture revenue from users without having to share it with the content owners. Even Microsoft's authentication and identification service, Passport, will be angling to control ownership of the consumer's information. In fact, the battle between AOL and Microsoft is more than anything about who controls the primary consumer interface. It is each company's hope that by gaining control of that interface, it will effectively have a platform that content and application providers will have to plug into in order to get access to consumers. The companies will, as a matter of course, charge for that access.

We assert that many new competitors will threaten to disintermediate end-user relationships currently owned by existing value-chain participants by disaggregating content from its historical distribution channels—but, to what extent? Given their current financial state, companies such as TiVo, Intertainer, and RealNetworks are just as likely to fail as succeed. And even more established "new era" companies, such as Yahoo and Gemstar, which have at least captured market value since their beginnings, have not yet demonstrated an ability to capture significant revenues in content- and media-related transactions. In other words, should existing players care?

For content companies, the basic question often has been: If I continue to create compelling content, isn't that enough? (In other words, no matter which distribution channel consumers choose, great content will be needed, and therefore I'll get paid my share.) The answer is: Not necessarily. While compelling content will continue to find an outlet, the increasing amount of user-generated content (e.g., chat, messaging) facilitated by digital technologies does reduce the share of overall content for which traditional content owners will be compensated. This is also at a time when the battle for a consumer's leisure time is already an issue; one need only look at the proliferation of magazine titles, Web sites, and television channels to appreciate the extent of the battle for the end consumer's attention.

In addition, it is less certain what share of the value of content license holders will capture, given the increasing number of potential intermediaries. It is likely that no single content company will be able to own all the content in a particular category, thus there will always be a need for an intermediary to aggregate that content. However, if content

FIGURE 3

Content and Media —The Digital Ecosystem

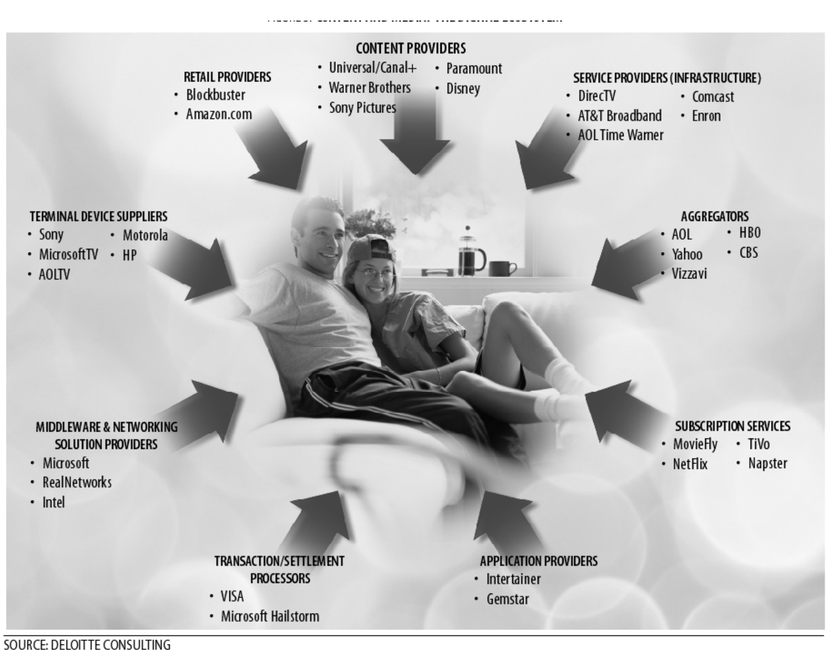

SOURCE: DELOITTE CONSULTING

companies cannot ultimately control the end-user relationship because of these new intermediaries, then the concern becomes how they will maintain and grow the share of the value that is captured vis à vis those intermediaries.

And, as production costs decline due to cheaper, digital alternatives, and as alternative, cheaper forms of digital distribution emerge, the core value that content companies offer to the content creators they represent (e.g., independent musicians and producers) needs to change. Today, the value that content companies offer to their content suppliers is in securing control over proprietary distribution channels. However, as those channels open up and the costs of production and distribution begin to come down, that value becomes more a function of marketing value, which is driven by the depth of knowledge content companies have of the end consumer's needs. Thus, knowledge about the end consumer will become the new dimension firms will compete along to attract and sign artists and producers.

For content aggregators—especially those dependent on advertising revenues—the proliferation of distribution channels enabled by new digital technologies, such as peer-to-peer and PVR, threatens to fragment the consumer base needed to drive advertiser interest. To illustrate this point, with PVR functionality enabling users to download programs independent of their original schedules, one could argue that TiVo (or DirecTV or AOLTV) becomes the broadcast network. Indeed, in the new digital ecosystem, the identity of your competitors is less clear.

In the broadcast (analog) world, networks were content aggregators that built brands to compete for ratings. In the broadband (digital) world, networks become the infrastructure that delivers content (bits) to consumers. As such, the greatest competition for networks may not be other networks, but instead the network operators through which the network programming passes (especially as these operators integrate TiVo-like functionality into their systems). In fact, those content aggregators that cannot control the distribution of their content or that do not possess a strong brand recognition with the end user may be replaced by their emerging digital counterparts. For aggregators, the threat of losing their relationship to the end consumer poses the even greater threat of losing their advertiser relationships, which account for the majority of their revenues.

For distribution companies, the growth of "open," Internet protocol (IP)–based distribution systems threatens to undermine their primary source of advantage—that is, their virtual monopoly over access to end users because of the incredibly high costs associated with duplicating their distribution infrastructures. Peer-to-peer distribution, for example, has the potential to be a "disruptive technology" that fundamentally changes the economics of content distribution. Peer-to-peer decentralizes distribution and is one of the ways content companies, content aggregators, and emerging digital distribution companies can open up what has been to date a closed-end system.

In such an open system, value will not necessarily flow to those with the greatest capital investment. Instead, companies may take advantage of the viral nature of the peer-to-peer medium (by offering, for example, innovative incentives to consumers for distributing and contributing content). These companies will focus on building their end-user relationships as opposed to their physical asset bases, and thus may change the way content is distributed. At a minimum, for cable and satellite companies—not to mention brick-and-mortar retailers—the possibility that wireline and wireless telephone operators and network resellers will siphon off their most profitable customers will force them to develop new ways of differentiating themselves to their consumer base.

All of this is not to suggest that the future content and media ecosystem will no longer consist of content creators, content aggregators, and distributors. Nor is it to suggest that scale no longer matters or that acquisitions will not take place—in fact, quite the contrary, given the recent indications from the Federal Communications Commission (FCC) that many existing ownership limitations will be relaxed. Rather, what will change is the relative value captured by each of the participants, as will the way that scale is defined. Physical scale will, to a certain extent, be necessary to maintain leverage in negotiations with ever-larger suppliers and buyers. But, with many of the larger media conglomerates approaching that size and with an explosion of content that only serves to exacerbate the conflict over a consumer's leisure time (the issue is not trying to find more content to push at consumers, but rather getting consumers to notice your content), then it follows that the nature of advantage will originate from other capabilities and assets.

"Traditionally," Thomas Middelhoff continued, "media companies said, 'We must control distribution.' In the future, corporations must establish direct access to consumers through multichannel on-line and off-line distribution strategies."

We believe that the move from analog to digital requires a new competitive framework, one that focuses on who owns the consumer relationship. And, we believe that the battle for owning that consumer relationship will not simply be through increased content and distribution asset ownership, but will instead occur along a third dimension: personalization.

Personalization: The Third Dimension of Competition

"Did you hear the one about the guy who was too poor to get personalized plates so he changed his name to J3L2404?" (Marge Gunderson [Frances McDormand], Fargo*)*

If traditional competition occurred along two dimensions, future competition in content and media will occur along a third dimension—a dimension we are calling personalization (see *Figure 4*).

What is it? Personalization is a measure of the depth of knowledge a company has about the preferences of its customers and that company's ability to deliver customized products and services to its customers based on that knowledge. While personalization is about having a broad range of customer contact points—across media, devices, and platforms—it is more than simply having a direct relationship to a consumer. Personalization may be enabled by owning physical assets, but it does not flow directly from their ownership. Nor is it simply a question of building a database with superior analytical tools to process that data, though that is a component. Companies need to develop technologies, business processes, and organizational structures to truly compete along this dimension. And they need to do so within the privacy guidelines that will be established in the coming years.

Why is it emerging? As described in the previous section, the digitization of content and the emergence of peer-to-peer file-sharing networks have created a new way for companies to compete against established content and distribution companies. The growth of alternative last-mile solutions and the adoption of open communications standards are changing the rules as well. While Napster's rapid acquisition of more than 50 million users can be explained to a great extent by its giving away content for free, one of the messages of its success should not be ignored—i.e., that new distribution mechanisms that deliver to consumers exactly what they want, when they want it, can dramatically impact existing distribution systems. In addition, with a preponderance of content alternatives and an ever-shrinking quantity of leisure time, consumers will favor companies that can simplify their lives and that are more attuned to their schedules and preferences.

What do content and media companies get from it? In a phrase: superior economic leverage. There is a geometric return associated with each additional user or each additional product sold to an existing consumer in companies that have successfully transformed themselves into personalization engines. Each time a user purchases an item on Amazon.com[1], for example, that information generates multiple, targeted marketing opportunities—not just to that user (e.g., notification when that author writes another book), but also to all other users with similar profiles (e.g., notifications based on their prior history that a new book may be of interest to them).

FIGURE 4

The Content and Media Competitive Landscape—The Three Dimensions of Competition in the Digital Age

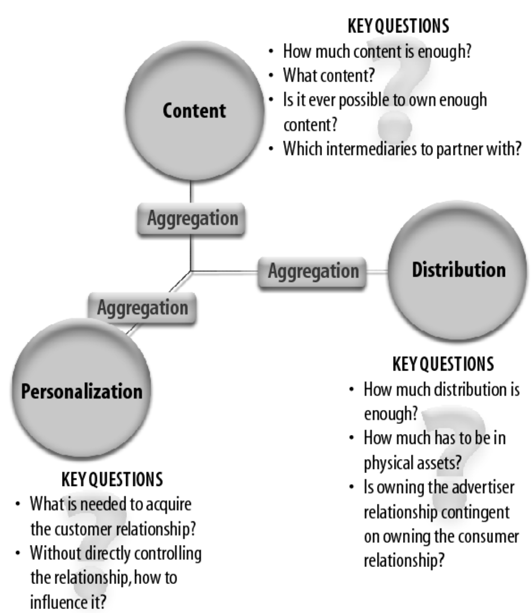

SOURCE: DELOITTE CONSULTING

In another example, Lands' End created a "My Personal Shopper" personalization engine on its site, and people who used that service bought on the site 80 percent more often, with an average order value that was 10 percent higher than orders made through other channels. And, in early VOD trials, monthly subscriber churn was reduced from more than three percent to just less than one percent. While it is still too early to tell what these numbers will be, as VOD becomes more common, the point that users stay with services that provide them what they want, when they want it, is an important one, especially as many content and media companies contemplate moving from single- (product) sale business models to subscription-based models that will be subject to churn.

In addition, personalization enables differential pricing for what is essentially the same product or service. This will allow companies to stress, for example, certain product attributes (e.g., timeliness or convenience) or even to charge different pricing models altogether (e.g., subscriptions, renting, buying, or bundling) based simply on a more sophisticated view of the consumer. The financial benefits of personalization include greater frequency of sales/visits, higher average price per sale, and lower customer acquisition costs as a percentage of sales, reduced churn. And, unlike building scale in physical dimensions, building scale along the personalization dimension can result in a much higher return on assets.

How can companies get it? Personalization requires more than technology—companies need to be completely oriented toward personalization to compete along this axis. In addition to the technology investments in relational databases, data-analysis tools, call-center applications, customer-relationship management (CRM) software, and so on, companies must also invest in the people and processes needed to gather and analyze information about user preferences. This is especially relevant for companies that have not strongly focused on the end user before, which is the case with many content and media companies that sell through distributors. Companies without direct access to consumer information will need to negotiate with those companies that do have access. While that is likely to be the most contentious part of the negotiations, the benefits achieved by delivering the right product at the right time should leave enough value for both parties to share.

How will companies use it to compete? In converging industries, where many companies, both old and new, are trying to position themselves to own the consumer relationship, consumer loyalty is up for grabs. Companies will use their personalization assets to build and strengthen their consumer relationships; this correspondingly raises switching costs for those consumers. As switching costs rise, the physical-scale equation changes—increased consumer leverage through personalization reduces the relative value imparted by physical assets, namely content and distribution assets.

For content companies, the threat of personalization is the risk of reduced leverage with other members of the value chain. Content aggregators and distributors will use their newfound leverage with consumers to negotiate more favorable terms with their content suppliers. As companies such as DirecTV, for example, know more and more about a particular user's preferences, that user may not be moved to switch even if a content company's channel is dropped or if a local channel is not carried.

For content aggregators, the threat of personalization is the risk of being cut out of the equation. We have already discussed how PVR technology threatens traditional advertising revenue, especially for programming that is not "live." This will force networks to pay ever-increasing rates for live programming, such as sports and news. To a great extent, we have already seen such a rise in license fees. However, if content aggregators do not develop a deep relationship with their audience, they run the risk of not being able to pay as much for those licenses as companies that do, even if those companies reach fewer people. Even if a distributor reaches only half the audience that a network might, it may see greater returns on the advertising it sells, and thus be able to pay more for licenses. This is because a distributor will be providing targeted advertising, which generates much higher costs-per-thousand (CPMs) than broadcast advertising. (And, in this example, reaching half the audience would require twice the CPM—a multiple that personalization has been shown to have.)

For distribution companies, the threat of personalization is the risk of losing their primary consumer relationship and stranding their significant physical asset base. For example, while existing cable and satellite companies may have an inherent advantage over their competition by virtue of their being the primary source for video programming because consumers are simply used to subscribing to their services, alternative sources of video subscription services that deliver personalized content may be able to draw certain consumers away. This is possible even if those services suffer, at least initially, from either reduced content (e.g., due to licenses being locked up) or greater inconvenience (e.g., due to users having to download programming to their PCs). By understanding what certain customer segments want and are willing to pay for, these alternative access companies can win share from those companies that treat—and charge—their customers the same. In other words, physical infrastructures built for mass-market distribution have the potential to lose out to lower-cost architectures tailored to specific customer segments.

What does it all mean? In the short term, content companies will be king, as many players, in an effort to build primary relationships with consumers, pay large sums for content. In the longer term, however, consumer-relationship managers will be king, and the most successful content and media companies will have a superior understanding of the consumer and the ability to deliver what she or he wants, to where she or he wants it, and when. For a discussion of who those relationship managers might be, please see the next section, "Strategic Uncertainty: The Possible Futures for Content and Media."

Why now (or haven't we heard this before)? Before moving on, however, it is worth discussing why companies need to be thinking about this today. For some time, we have heard people talking about how there would not be 500 channels, but rather only one ("my channel"). And, the concept of mass customization, while not synonymous with personalization, gained favor in the early 1990s, only to fall out of favor as companies realized that the costs of customization far outweighed the benefits.

Indeed, personalization is not free. There are always costs associated with implementing it, not the least of which is the loss of basic economies of scale that are achieved through mass production. In the case of digital music, on-line radio broadcasters need to pay additional royalties for allowing consumers to develop personalized play lists, the theory being that a more personal level of connection with the consumer will eventually lead to additional revenues from that consumer, improved word of mouth, and other benefits. Personalization does have to provide real value that makes enough of a difference to the consumer, and that is a lot more difficult—and expensive—than a lot of people thought it would be.

So, what has changed? First, technological developments have made personalization more feasible. The PVR, for instance, allows for the simple creation of a single, personalized channel, and advances in CRM systems provide the needed infrastructure to support personalized service offerings.

Second, during the past decade, media have proliferated dramatically—there are thousands of new magazines (more than 16,000 consumer titles), hundreds of new television channels, and thousands of Web sites—causing an even greater need for a company to filter out content that is not relevant to a consumer.

Third, there has been a change in consumer behavior itself. A new generation of consumers that were raised on interaction with content providers and personalization of their Web sites demonstrates consumers' increased comfort level with personalized services.

Fourth, much of the growth in the past decade was achieved through acquisitions. Many of the content and media companies, in their current form, are experiencing slower growth—cable penetration has, for the most part, peaked; dial-up Internet access is slowing down; domestic music sales have actually declined; the number of movie tickets sold is fairly flat—and there is a question about just how much more these companies can use acquisitions, at least domestically, to grow their markets. What's clear is that companies need to develop new ways to find growth.

Fifth, the AOL and Time Warner merger speaks to a new type of company, one focused on the subscriber and with interactive/personalization capabilities at its core.

While all of this suggests some fundamental changes in the way companies will be competing going forward, it is worth making the point that competing along the personalization axis requires an explicit, well-executed strategy. Past attempts at "mass customization" failed not so much because the basic concept was flawed, but because companies did not analyze what type of customization was valued by their customers. Personalization for personalization's sake will indeed lead to higher costs without any associated benefits. Companies need to pursue the type and level of personalization that is precisely right for their customers, and this can only be gauged through thorough and ongoing customer analysis.

Strategic Uncertainty: The Possible Futures for Content and Media Companies

"I don't know the future. I didn't come to tell you how this is going to end. I came here to tell you how it's going to begin." (Neo [Keanu Reeves], The Matrix*)*

While we believe personalization will play a major role in the future, the tremendous amount of regulatory, technological, and end-user uncertainty makes it unclear who the leading consumer-relationship managers will be. We have defined three scenarios that, in our opinion, describe the possible future evolutions of the content and media businesses. As part of any such exercise, we will attempt to describe the underlying factors and events that need to occur to make that scenario happen (see *Figure 5*). After describing each scenario, we will discuss the possible strategies that different content and media companies might use to influence the evolution to their advantage.

1. My Studio

This scenario describes a world where content companies and content aggregators—and their personalization assets—capture the majority of the value. Here, the stranglehold that existing closed-end distribution systems have over consumer access is broken. This implies that the deployment of alternative last-mile solutions into the home occurs, that open standards are developed to tie the various components in the digital living room together, and that regulatory agencies push for accelerated broadband deployment as well as for must-carry provisions from existing distribution companies. Given these events, distribution assets become increasingly commoditized, and content companies have the opportunity to develop more direct relationships with the end consumer.

2. My Walled Garden

This scenario describes a world in which closed distribution systems—and their personalization assets—capture the majority of the value. In other words, "proprietary" broadband access rules the day, given the slow deployment of alternative distribution channels due to inferior economics and difficult-to-change end-user behavior, a passive regulatory regime, and minimal standardization of technologies. In this scenario, assuming the distribution companies simultaneously leverage their position to strengthen their end-user relationships, certain content companies and aggregators may be marginalized and suffer in terms of negotiating power versus these distributors. And, as more and more functionality—including, for instance, PVR and messaging—gets built into these closed-end networks, the role of traditional content aggregators that do not have either a strong brand or access to unique content will be threatened.

3. My Entertainment Services Provider

This scenario describes a world in which pure relationship managers capture the majority of the value. In other words, companies that have amassed the largest database of consumer behavior and built the capabilities to leverage that data into personalized content and media offerings will not only exist, but will capture greater value than pure content and distribution companies. (Of course, the question of how these companies acquire those end-user relationships without the physical assets to begin with merits discussion. The way Amazon.com has positioned itself and the way

FIGURE 5

Strategic Uncertainty—The Possible Futures for Content and Media

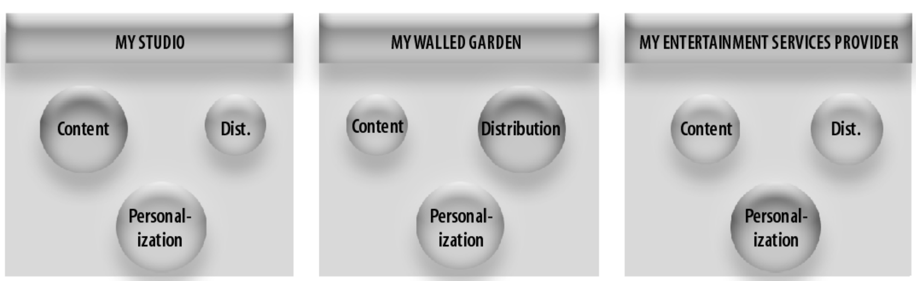

MY STUDIO	MY WALLED GARDEN	MY ENTERTAINMENT SERVICES PROVIDER
FACTORS	**FACTORS**	**FACTORS**
• Last-mile solutions commoditize access • Open standards • Must-carry	• Proprietary broadband access rules the day • Minimal standards (wild west) • Passive regulatory	• Broadband ubiquity (channel proliferation) • Device convergence • Must license
POSSIBLE EVENTS	**POSSIBLE EVENTS**	**POSSIBLE EVENTS**
• Vivendi/Yahoo • Viacom/Amazon • Sony/RealNetworks	• DirecTV/TiVo/Blockbuster • Comcast/AT&T BB/Electronic Arts	• Yahoo/MGM • Microsoft/AT&T Broadband
STRATEGIES	**STRATEGIES**	**STRATEGIES**
• Amass content breadth • Acquire end user relationships • Acquire personalization technologies and platforms • Establish brand • Sell packages versus products	• Acquire limited content • Consolidate closed-end systems • Acquire personalization technologies and platforms • Build customer loyalty as standards develop	• Acquire limited content and/or distribution • Amass customer interactions • Establish brand • Virtual integration • ESP where necessary

SOURCE: DELOITTE CONSULTING

RealNetworks is positioning itself suggest the opportunity for software-based—as opposed to physical asset–based—companies to play a significant role in the value chain. In addition, one can also assume that a distribution or content company, once it has developed a sufficient number of high-quality end-user relationships can, over time, divest their physical assets to focus on their highest performing, lowest capital-intensive assets.) In situations where the entertainment service provider cannot develop that end-user relationship, it can make its engine available to third parties (via a private-label arrangement) but still participate in the revenue, especially as its broader database facilitates up-selling and cross-selling for its partners.

This scenario requires broadband ubiquity enabled by alternative access technologies; channel, content, and device proliferation; and "must carry" and "must license" provisions from regulatory bodies to ensure supply from distribution and content companies. While this scenario is not likely to occur in the immediate future, the very real possibility of access commoditization as well as the reality that no single content company will ever be able to assemble all of the content under its umbrella do suggest the possibility of a future where a standalone entertainment services provider will emerge.

It is important to note that the likelihood of each scenario actually occurring will vary by content type and that, in fact, what is likely to happen will be a fragmentation in the evolution of different content and media businesses. For example, while video content looks to be delivered through the "walled garden" of existing distribution systems (at least in the short to medium term), music content is already finding alternative distribution outlets that suggest different outcomes. The scenario-planning exercise described needs to be more rigorously applied to specific content categories, though it does point to certain lessons that will apply across categories and suggests the issues that content and media companies need to start considering in order to control their futures.

That said, what are the strategic implications for the various content and media companies, and where should they be focusing their efforts to steer the evolution to their advantage? Given the number of new companies entering the digital content and media space, the rapid degree of technological change, and the highly uncertain response of end consumers—i.e., given the likelihood that the different scenarios each have a probability of occurring—highly flexible strategic planning will be required. That means, for instance, the use of real (strategic) options to hedge and the constant monitoring of when those strategic options are in or out of the money. Often, it also means going down several parallel paths—each of which has implications for what the optimal organizational structures and buy-versus-ally strategies might be.

While strategic flexibility will be the overarching guiding principle for these companies—at least for the next several years—there are specific actions that each segment of the content and media value chain can begin taking.

For content companies and content aggregators, the principal strategic challenge will be to commoditize distribution assets or, at a minimum, to maintain leverage with the distribution companies. That means amassing a breadth of unique content and investing in their brands, though the level at which branding occurs—at the individual program level or the network level, for example—will be an interesting issue for content companies. If content and media companies do see a world of open delivery systems, then investing in distribution assets may not be the best deployment of capital, which can be better invested in brand and content development.

Leverage with distribution companies can also take the form of allying with or even acquiring major suppliers to the distributors. News Corp.'s ownership of Gemstar, for example, is a point of leverage it can use with domestic cable and satellite operators despite its relatively light position in domestic distribution assets (although a successful acquisition of DirecTV will certainly change that). Pushing open standards is another mechanism to combat the power of distribution companies, although the question of which standards is an issue. Peer-to-peer standards as well as open standards for set-top box manufacturing—to push them into retail outlets—are ways to open the door into closed-end distribution systems.

In terms of moving along the personalization axis, content companies will need to start investing in personalization technologies and platforms, if for no other reason than to take an option on their future success. AOL Time Warner's recent investment in Amazon.com is a step in that direction, as is Vivendi's acquisition of MP3.com. Other possible events include partnerships with and/or acquisitions of Yahoo! and RealNetworks, with Vivendi and Sony as the potentially interested parties, given their relative lack of presence in networks and branded content aggregation assets.

While content creators (e.g., film studios, record companies) and content aggregators (e.g., networks, cable channels) face similar challenges, they each face additional challenges unique to their position in the value chain. For content creators, a major challenge will be how to develop direct relationships with the end consumer, given that aggregators and intermediaries (e.g., bookstores, video stores, movie theaters—and their digital equivalents) will continue to play a role in offering a breadth of selection from content companies.

The additional challenge, as we have discussed, is that digital distribution for the first time divorces the content from the distribution channel, a pairing that content companies use today to generate most of their money (e.g., CDs sold as units in retail stores). Where content companies can establish those direct relationships, they should—although managing channel conflict will be an issue. Where they cannot, the key will be making sure that their content does not get divorced from the complementary revenue streams associated with that content.

In other words, instead of pushing units to their distribution partners and not receiving the end-user information, content companies need to start thinking about the packages of content they can assemble and then distribute. For example, music companies need to sell more than just a CD (or its digital equivalent) into the channel. With the core music file that gets distributed, a music company could also package-in links to its Web site, an interactive game, a chat session with the musicians, a concert-tracking and ticket-purchasing application, and various sweepstake and marketing offers, all of which enable the content company to capture end-user information and to pre-empt content distribution companies from capturing the value of complementary products and services.

For content aggregators, on the other hand, there is the opportunity to become an entertainment services provider, managing consumer profiles to deliver personalized content across the disparate closed-end distribution systems—i.e., to the device and through the channel specified by the consumer for that particular category of content. Given device proliferation, as well as the increasing number of services that consumers are expected to subscribe to, there is an opportunity for a company to aggregate those services for the end user, thereby utilizing superior personalization assets to bypass the existing advantage of closed-end distribution assets.

For existing content distributors (namely, cable and satellite companies), the consumer relationship, to a certain extent, is theirs to lose (although cable and satellite companies will continue to battle for market-share).

In the short term, content distribution companies will continue to consolidate their positions in closed-end systems through acquisitions as well as through the development of proprietary standards. In addition, these companies will acquire limited content in order to maintain a position of leverage over their content suppliers. DirecTV's growth was tied to its purchase of exclusive rights to sports programming, and one can envision a similar acquisition to initiate control over interactive media applications.

DirecTV or Echostar might buy a company such as Electronic Arts, given that games distribution will increasingly be through digital channels, and then develop proprietary gaming applications that encourage users to sign up with them. At some point, owning a few prized content assets will enable these companies, with their direct user relationships, to improve their position vis à vis standalone content providers.

In the longer term, the challenge for these companies will be to move from being passive deliverers of billing statements to active managers of their consumer relationships. Or, to put it another way, a satellite company such as DirecTV needs to think less about being a "satellite company" and more about being an "entertainment services provider," delivering content to consumers based on their unique insight into their consumers' preferences. This may actually mean managing the delivery of content over systems not necessarily owned by the company, such as delivery over wireless networks to cell phones or even IP networks to personal computers (PCs). As part of this transformation, these companies will need to invest in personalization technologies and platforms, with DirecTV's investment in TiVo being one example.

We may even see mergers between brick-and-mortar companies—companies that, despite being very much at risk from digital distribution technologies, possess very strong consumer databases —and their digital equivalents. Blockbuster's partnership with DirecTV to sell subscriptions is an example of the potential mutual benefits stemming from this type of relationship, though the real benefit to a company such as DirecTV might be in Blockbuster's infrastructure for managing consumer information and for marketing to consumers.

For new content distributors such as wireline and wireless telephone companies, there is the double challenge of penetrating the market owned by the existing distributors while simultaneously avoiding the commoditization of their infrastructures to common carrier status by content companies.

One of the biggest issues here is educating consumers that there are alternatives for them when it comes to content and media suppliers. The most probable strategy for these companies will be to pursue one or two specific content and media categories that lend themselves to these alternative distribution channels. While there is much talk about IP–delivered VOD services, they most likely comprise one of the categories that will not be delivered through IP in the near term. Games and music are the most likely categories for new content distributors to pursue.

Like their cable and satellite brethren, these companies may need to make a targeted content acquisition or partnership in order to drive interest in their services. Once consumers begin changing their behavior with regard to the consumption of entertainment, then these companies can start bundling in additional content categories.

The major advantage new content distributors will have over their competitors is the ability to build more flexible architectures than existing distributors' mass-market, broadcast infrastructures. So, while a full array of video programming may be difficult to assemble for these companies, they may be able to target the segments of customers that would prefer a lower monthly television bill because they only watch a very small and focused portion of the programming.

Like the content companies, new content distributors need to find ways to open up closed-end distribution systems. Open standards are one way of accomplishing this. It is our opinion that peer-to-peer technologies give new content distributors (as well as the content companies themselves) the ability to fundamentally alter the economics and dynamics of existing content distribution. And it is, potentially, the new content distributors (wireless and wireline companies) that can lend legitimacy to the peer-to-peer movement.

Owning the Consumer Relationship: Challenges for the 21st Century Content and Media Company

"Are you ready to surf, Lance?" (Lt. Colonel Bill Kilgore [Robert Duvall], Apocalypse Now*)*

While the possible futures are uncertain and while the ways in which companies plan for the future world of personalization will be different, what is certain is that content and media companies will need to develop competencies in personalization that cut across all aspects of their organization. As discussed earlier, personalization requires that companies undergo a complete reorientation and rethinking of the way they do business.

Distribution companies need to move from thinking of themselves as simply pipe companies sending bills to their customers and more as the managers of their individual subscribers' entertainment experiences. Content aggregators need to define, in very explicit terms, what their brand means to their audience, or the new networks will steal their audience. And content owners need to think less about selling products into the distribution channel and more about delivering cross-category experiences, both indirectly and directly, to the end consumer.

The transformation will be a wholesale one for many companies—in particular, for those that have very little interaction with the end consumer today—thus affecting all aspects of the organization. The following are the areas of consideration—and questions to ask—for companies looking to own the consumer relationship.

- ***Strategic Scope:*** The nature of scale will change as the third dimension of personalization is added to the competitive equation, as will the extent to which companies need to be vertically integrated to compete.
  - What assets are needed in order to compete along the personalization axis?
  - How many assets are still needed along the content and distribution axes to ensure adequate access to consumers and to information flows about their behavior?
  - At what point can physical assets be shed (i.e., at what point do physical assets have their highest potential value but lowest future value)?
- ***Customer Strategy:*** Competition will increasingly take place at the segment—even individual—level, and customer strategy will need to follow.
  - Which customer segments are most vulnerable to companies offering personalized services?
  - How are customer segments even defined in industries undergoing convergence (e.g., the digital living room)?
  - What offers should be made to those newly defined customer segments?
- ***Pricing Strategy:*** Many companies will move to dynamic, real-time pricing models based on specific consumer behavior in order to differentiate themselves.
  - At what level are pricing decisions now made (e.g., at the individual or segment level)?
  - How do you move to a dynamic pricing model, where pricing occurs in real time (or near real time)?
  - How do you get dynamic pricing accepted by consumers?
- ***Organizational Structure:*** Companies will need to migrate from product-based silos to flexible structures organized around the consumer.
  - What is the appropriate balance between centralized and decentralized control in companies competing to own the consumer?
  - How do compensation systems need to change in order to reward cross-selling and sharing of consumer information?
- ***Financial Systems:*** Companies will need to move from tracking individual properties/products in their financial systems to tracking individual consumers.
  - What is the lifetime value of a customer?
- ***Billing Systems:*** Companies whose billing is based on unit-based pricing will need to migrate to systems based more on usage- and subscription-based pricing. Companies whose billing is at the household level will need to consider billing at the individual level.
  - What part of billing (if any) can be outsourced, given this fundamental change?
- ***Customer Care:*** Companies trying to build direct consumer relationships will need to build supporting customer-care infrastructures that are integrated across the various customer contact points (Web, call-center, retail, etc.).
  - To what extent can the customer-care infrastructure be leveraged into sales and marketing functions?
- ***Technology Architecture and Other Systems:*** Companies will orient their systems around the individual. Database architectures that are not relational in nature will need to become so. Digital asset management systems need to be less about managing units and more about assembling packages.
  - What architecture protects the company's assets, ensures consumer privacy, and flexibly responds to technological change?
  - What does rights management mean when the Internet transcends traditional territory definitions?
  - What does digital asset management mean, as content companies move from products to packages?

The biggest challenge for companies will be making these wholesale changes as they continue to operate their businesses, especially because many of these new investments will grow at the expense of some of these companies' existing revenue streams. For example, maintaining the satisfaction levels of existing channel partners while pursuing direct consumer relationships will be a major challenge—and one that some companies will simply choose not to assume.

Internally, balancing the rewards for individual group performance against the rewards for an integrated end consumer focus will prove equally challenging. And, of course, if the rate of transformation exceeds the rate at which these new markets develop, then overall shareholder performance will decline.

Indeed, these are not trivial challenges. However, they are challenges that must be met to compete along this third dimension of personalization and to win in this new, digital era.

Are you ready to surf the personalization wave?

Notes

1. *A note on Amazon.com:* While it certainly could be argued that Amazon.com does not have the financial success to be used to support this point, we feel that the company's financial difficulties are less a function of a flawed personalization strategy than a function of an overly aggressive content strategy—i.e., offering too many categories, often in areas that have very different logistics and support costs from the core offerings. Staying focused on books, videos, music, and video games—which all contain similar distribution logistics—and then licensing the personalization engine to other retailers of other merchandise would have enabled Amazon.com to capture the advantages of the data it collected, without all the associated costs. From that base in content and media-related products, Amazon.com could have used its resources to move into video rentals, music-on-demand subscription services, and the like—all powered by its strong personalization capabilities—and ultimately shown better financial performance.

Summary of Strategic Trends in the U.S. Telecommunications Industry

Joseph S. Kraemer
Director
Law and Economics Consulting Group (LECG)

Introduction

Strategic Trends presents perspectives shared by a wide range of industry participants on what has occurred, is occurring, and, most importantly, can be expected to occur in the telecommunications industry. This paper has two objectives: (1) bring some degree of consensus out of the confusing strategic situation of late 2001 and (2) provide input to the strategic planning processes of companies struggling to survive the current downturn and then thrive during the recovery.

Strategic Trends constitutes a joint effort. More than 50 senior industry executives and providers of professional services (e.g., bankers, attorneys, consultants) contributed to the final version.

The method employed was a modified Delphi in which industry experts react to a hypothetical "future" with a mix of agreement, criticism, and alternative explanations and interpretations. Opinions were not attributable to individuals and represent personal views, not "official" corporate positions. Given that the participating experts have decision authority in their organizations, it is probable that they will act in accordance with their perceptions/beliefs, thereby creating the classic "self-fulfilling prophecy." Therefore, whether you as an individual agree or disagree with the consensus, you can expect that the strategic and tactical plans of participating industry stakeholders will reflect the consensus view of *Strategic Trends* or a close variation thereof.

Strategic Trends began as a late August 2001 e-mail brainstorming session among three consulting firms. Clients of those firms were then asked to comment. More and more input has occurred as versions of the document have been forwarded throughout the industry, posted on Web sites, and presented at industry forums.

Within two weeks of the tragedy, decision-makers began to react to the events of September 11, 2001, and its aftermath. Several issue areas in this version reflect post–September 11 abandonment/modification/reinforcement of previously held views. Such changes in perspective reflect both (1) the commercial consequences to telecommunications companies of "exogenous" economic/political/social events that originate outside the industry and (2) the fragility of the current economic outlook that is vulnerable to such unanticipated events.

The paper is organized as a series of issue areas. These issues are those that are most significant to the future of the industry in 2002 and beyond.

Recovery No Sooner Than 2003

In the <u>best case</u> scenario there will be no material recovery in the telecommunications industry until at least the summer of 2003.[1] The delay in recovery means that the slump will be longer and deeper than the late 2001 conventional wisdom.

The telecommunications industry will *lag*, *not lead*, the recovery of the U.S. economy. The lag will be two to three quarters (i.e., for the industry to recover in mid-2003, the economy must experience a sustained recovery before the end of 2002). The period from now until the recovery will be one of "triage" during which dead and dying businesses will be identified and closed out so as to maximize the remaining opportunities for the survivors. The triage process began in late 2000 and will continue at least through 2002.

Aggressive and well-managed companies will recover sooner. These companies will take advantage of the industry situation to improve their competitive position through both internal growth and acquisitions. In effect, aggressive managers will out-hustle and out-maneuver weaker, less-well-managed competitors.

"Supply Push" dominated the 1996–2000 period, when aggregate network capacity came on-line far in excess of near-term demand. The next growth period (2003 [?] and beyond) will be characterized increasingly by "Demand Pull" as (1) customers sort out the applications that they require and the amounts that they will pay and as (2) younger and more "wired" generations (born 1970 and after) start households and increase their discretionary income. *The management challenge is to take maximum advantage of the opportunities present as the industry consolidates and retrenches* (see *Figure 1*).

FIGURE 1

Back-to-Back Growth Curves

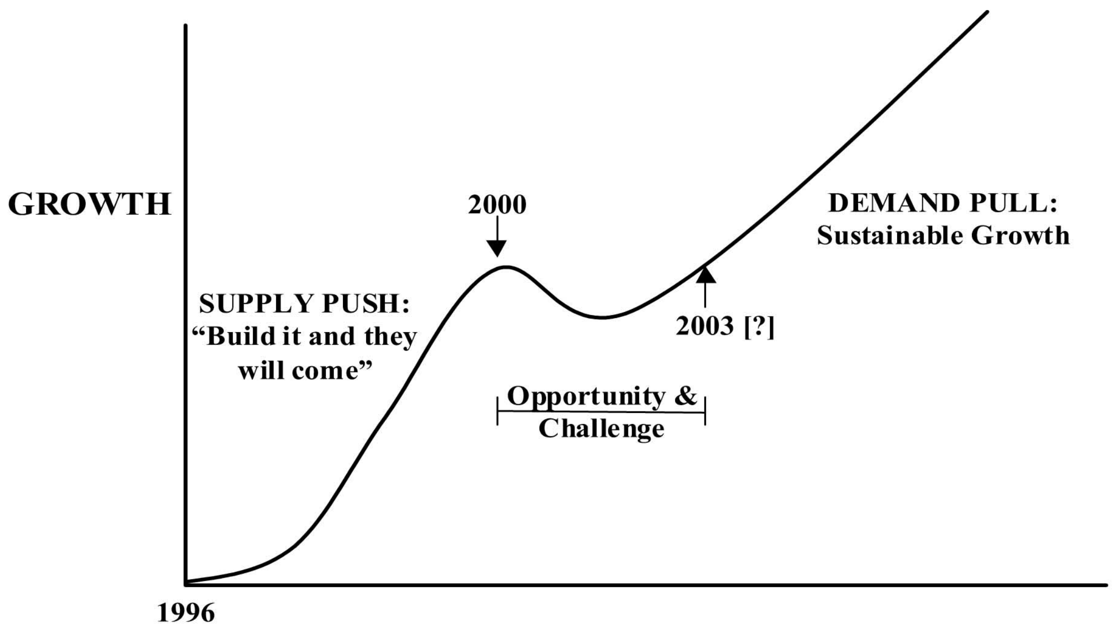

The Future Consolidation of the Industry

The carrier segment of the telecommunications industry is consolidating with an eventual (circa 2004–2005) configuration similar to the airline industry—three or four majors, a few regionals, and some niche players. Accordingly, the distinction between "local" and "long distance" will disappear. Telephone service (including local and long distance) will be delivered by vertically integrated suppliers.

The total number of licensed wireline carriers in the United States will decrease by 75 percent between year-end 2000 and year-end 2003. In contrast to the chaos of the last 18 months, markets will stabilize. There will be many fewer carriers to offer price, product, or service competition. However, customers will have confidence that their carriers of choice will remain in business and deliver services in accordance with contracts (i.e., be viable in the long term).

The September 11 attacks reinforce the need for robust, interconnected networks that have a high probability of survival in the event of natural or man-made disaster. That argues for a consolidated base of carriers operating with agreed-upon disaster protocols.

Financially stressed carriers will have a debt load that is a burden. Therefore, acquisition of such a carrier (on an asset value basis) may be financially viable only after the carrier files for bankruptcy or restructures in a way so as to disadvantage existing debt and equity holders.

Consolidation will not be easy, particularly for smaller carriers. If one carrier acquires another, then the result is more customers on separate networks that may not be technically capable of integration at the physical and logical layers. In the short run, this works against achieving economies of scale. In most cases, the post-transaction period continues for 12 to 24 months. During that time, the risk is that the acquirer's management loses market focus and becomes distracted by organizational design and personnel issues.

Consolidation in North America will increase the pressure to create continental-sized carriers in Europe. Such European consolidations are more likely to receive European Commission (E.C.) regulatory approval subsequent to the creation of mega-carriers in North America.

Carriers will attempt diversification at the same time as consolidation. Carriers will stress investment and "strategic partnerships" outside their core businesses. They will continue to experiment with diversification into content and e-commerce applications. However, although some will be very successful, most will fail. Diversification is a high-risk strategy principally because the distinctive competencies associated with the content and e-commerce businesses are so different from the skill set that makes carriers successful.

The RBOCs as Survivors

Of all the participants across all of the segments in the industry, the regional Bell operating companies (RBOCs) are the most likely to survive and thrive. In fact, the most visible sign of the industry recovery will be an improvement in RBOC financial performance.

The RBOCs will dominate North America and most likely absorb *at least* the traditional long-distance businesses of AT&T, WorldCom, and Sprint. The local infrastructure developed over the last decade by AT&T and MCI would also be included.[2]

RBOC executives and other observers forecast that RBOC versus RBOC competition will accelerate and intensify as (1) industry consolidation takes place and as (2) the RBOCs transform themselves into national and global full-service providers. *This is a very controversial position with which non–RBOC commentators disagree*. Instead, the dissenters argue that there are no economic incentives for inter–RBOC competition. Therefore, these dissenters argue further that the RBOCs will structure their regions as "fortress hubs" and not compete in their core businesses or geographies (known as the "polite oligopoly" scenario).

In addition to inter–RBOC rivalry, the primary sources of competition for the RBOCs will be (1) surviving competitive local-exchange carriers (CLECs) operating as regional or niche service providers, (2) mobile wireless companies, and (3) cable multiple-system operators (MSOs) in residential and small-business markets. The RBOCs remain vulnerable to a video/Internet access/telephony service bundle sold by the MSOs.[3] However, a critical unknown is the extent to which the MSOs can and will finance and manage the deployment, marketing, sale, and provisioning of such a service bundle.[4]

The Evolution of the CLECs

The first phase of the CLEC era ended with a bang that is still reverberating through the capital markets. Lost in the noise of so many failures is the fact that a strong demand remains for an alternative local service provider to the incumbent LEC.[5] *Surviving CLECs will be predominately facilities-based[6] and serve small and medium-sized business customers. The survivors will function as regional or niche providers, not as "integrated communications providers" on a national basis.*

To be viable CLECs must (1) have access to operating capital, (2) achieve a minimum level of scale economics, (3) develop a high-loyalty/low-churn customer base, and (4) follow a very focused business plan. One viable CLEC model is an independent telephone company extending service into an adjacent RBOC's service area.

The first phase of the CLEC industry was characterized by (1) too much cheap capital, (2) very bad business plans, (3) poor management, (4) RBOC unpreparedness, and (5) inappropriate regulatory theories and actions. The most visible legacies of the first phase are bankruptcies, litigation, and the collapse of investor community confidence.[7] Regulatory/legislative intervention to "save the CLECs" (e.g., divide the RBOCs into separate wholesale and retail businesses) will distract management (of both the RBOCs and the CLECs), confuse investors, and, potentially, become a major cause of delay in overall industry recovery.

Suppliers (e.g., network equipment manufacturers and professional services firms) built unneeded capacity to serve the collapsing CLEC industry. The suppliers themselves must now, in turn, consolidate and shrink. Recovery for suppliers will *lag* recovery for carriers by one year or more (i.e., if carriers recover in mid-2003, suppliers will recover in mid-2004).

Broadband Progress

High-speed Internet access penetration of U.S. households may have almost doubled in 2001. However, final year-end data will show that high-speed Internet access adoption by residential and small-business customers fell below the original 2001 forecast.

The original industry consensus forecast was 11 million residential and small-business connections by year-end 2001.[8] Consensus estimates are between 9.0 and 10.0 million as of December 31, 2001. The expectation is that the growth for 2002 and 2003 will also be below industry expectations. The primary causes of the shortfall appear to be (1) price sensitivity under conditions of economic uncertainty and (2) a lack of ubiquitous digital subscriber line (DSL)–based solutions in urban and suburban areas.

Analysts distinguish "availability" from "take rate."[9] Cable and telephone companies in combination claim that high-speed access *is available* in urban and suburban areas and that the *take-rate shortfall* results from a lack of content and applications that appeal to consumers and small businesses. Conversely, companies in the content/applications businesses argue that high-speed access availability is either overstated or overpriced, or both.

The near-term demand for high-speed access will be driven by existing content and applications. Over the longer-term, new content and applications will be created as one of the results of the increased base of potential customers with high-speed access. *The synergistic effect is that of a "virtuous circle," with access and content/applications reinforcing each other in a positive direction.*

The pace of high-speed Internet access rollout is more than an esoteric discussion, because the conventional wisdom is that acceleration of high-speed access will affect the economy favorably. The stimulus will occur because of (1) the spending stimulus associated with the capital expenditures to build the networks and (2) the benefits of expanded business-to-consumer (B2C) electronic commerce. Beneficiaries of broadband include the local network companies (cable operators as well as LECs), along with a myriad of others including network equipment vendors, personal computer (PC) makers, chip manufacturers, software companies, and media companies such as America Online (AOL)–Time Warner. In addition, there are thousands of companies for which the Internet is a low-cost distribution channel. Secondary effects would be experienced by the suppliers to these companies.

Regulators and policy makers debate increasingly whether slow adoption is (1) a market failure that requires intervention, (2) typical at this point in time in the rollout of a potential mass-market service, or (3) a situation that requires a federal "industrial policy" regardless of whether or not the market has failed. The outcome of the debate will affect the nature and extent of government intervention.[10]

The current dispute as to whether RBOCs must unbundle newly constructed fiber networks for sale to competitors will slow down the near-term deployment of telephone network–based high-speed access capabilities. This will have two effects:

(1) Cable-based systems will become the method of choice for high-speed access.
(2) Deployment of fiber by the telephone companies to the neighborhood, curb, or home will be delayed materially.

In the event that terrorist activities force quarantine restrictions on personnel movements due to biological or chemical attack, dispersed work sites (including homes) with high-speed access will be required.

The Critical Nature of Video

The ultimate bandwidth-eating "killer application" will be video. Video applications will include the following:

(1) Entertainment video downloaded from the Internet as well as received over cable or satellite system
(2) Business video, including videoconferencing (potentially from nontraditional work sites such as homes) and corporate training (real-time class room and user-initiated downloads)
(3) Distance learning for formal education (K-12th grade, post secondary, and advanced) and personal improvement
(4) Web-based video content used for everything from selling automobiles to on-line classified advertising
(5) Video attachments to personal and corporate messages (including "instant messaging")
(6) Video telephony, which has a high approval rating from upscale demographic segments born after 1970

Video compression techniques will continue to improve, which will reduce transmission bandwidth for a given task. However, the net effect of compression techniques versus application growth will still require substantial low-cost bandwidth. Such bandwidth will come from "lighting" the unlit fiber-based systems now in excess supply. These fiber systems have been, or will be, sold to carriers that will buy and book them at a cost basis that will provide a positive return on investment (ROI).

The increasing prevalence of video-based content constitutes a technology and market-driven discontinuity that facilitates entry by nontraditional companies, such as AOL and Microsoft, into the communications business. In a video world, the traditional carriers are disadvantaged relative to new entrants. To the extent the embedded network infrastructure is not video-capable, then it becomes transformed from an asset to a liability.

Transport technologies will include wired and wireless systems, the latter being both terrestrial and satellite-based. Carriers such as incumbent local-exchange carriers (ILECs) that are not video-capable will operate at a competitive disadvantage.

Slower Than Expected Wireless Data

The U.S. adoption curve for wireless data services will be much slower than the experience with DoCoMo's wireless Internet service in Japan. The situation in the United States is complicated by extreme uncertainty as to the value and availability of wireless data spectrum.

Desirable spectrum is occupied currently by a combination of (1) television broadcasters, (2) federal government agencies (including the Department of Defense [DOD]), and (3) local school districts. All three present difficult technical and political issues that must be resolved prior to spectrum becoming available.[11]

Adoption and implementation of wireless data applications by businesses are accelerating. Business use of wireless data lacks the visibility it deserves due to media focus on consumer applications. However, over the long term, business data applications will be limited unless additional spectrum becomes available.

Wireless carriers must choose between alternative wireless data business models: (1) the *Traditional Model*, in which carriers provide an open platform and receive revenue primarily from traffic carried and some value added services (e.g., billing) or (2) the *Transactions Model*, in which carriers share in the revenue stream of content providers (i.e., the equivalent of an LEC sharing the revenue of Yahoo! or AOL on the wired network). The risk of the Traditional Model is that the revenue stream, gross margin, and earning before interest, taxes, depreciation, and amortization (EBITDA) may not generate a sufficiently robust ROI to justify the use of scarce capital. The risk of the Transactions Model is that carrier revenue demands and proprietary requirements will delay or stifle the development of the applications necessary to drive the wireless data business.

Wireless carriers may elect to share network facilities. This will reduce up-front capital requirements but restrict downstream competition to price (rather than technical features or coverage).

Network Construction Curtailed

The "surge era"[12] in carrier network construction has ended. The "network" in the United States (as completed or funded *and* under construction currently) is more than adequate to meet demand forecasts through the middle of this decade. The deployment of next-generation carrier network technology will be delayed until the existing surplus of current generation technology is absorbed.[13]

CAPEX will be curtailed back to approximate pre-1996 levels (on an inflation-adjusted basis). "Build it, and they will come" has gone.

The drop is most obvious in inter-city networks that are overbuilt. Construction will continue for intra-city fiber networks, but only for customer-specific builds with assured demand under signed contract.[14] The focus now will be on achieving the carrier's ROI that was contained in the business plans that justified the original network investment. (This may not be easy, since many of these business plans relied upon excessively optimistic scenarios and associated financial forecasts.)

There will be substantial emphasis during the 2003–2005 period on rolling out "home networks" that are essentially local-area networks (LANs) for upper-income households.[15] Such networks will be video-capable and are expected to increase demand for "last-mile" high-speed access service from cable and/or telephone companies.

Prices for transport and switching will continue to move toward the *cost* of the most efficient provider. Public policy will support this price movement despite the economic stress on many carriers with high embedded costs.

Extreme uncertainty exists currently as to baseline/equilibrium prices for telecommunications services. Short-term contracts tend to be the rule currently, since large buyers are avoiding long-term commitments. The short-term (next six to 12 months) combination of (1) excess supply and (2) repressed/tentative demand reinforces mid-2003 as the earliest recovery possible. Consequently, carriers will emphasize efficient, low-cost operations and scale economies that, in combination, will generate high margins and positive EBITDA.

Economics dictate migration onto a single packet network for voice, data, image, and video. In the long-term, maintaining two parallel networks—circuit-switched for voice and packet-switched for data—is excessively expensive.

Internet protocol (IP) telephony is more economical due to lower hardware/software costs and the greater efficiency of packet transmission. Combined capital and operating cost savings are estimated to be in the range of 30 percent to 50 percent. However, given the embedded network infrastructure, in the short run, many carriers plan to use Class-5 circuit switches for voice and to offload data to packet networks.

Large businesses will build packet-based extranets and intranets in which voice will be integrated and quality of service (QoS) assured. In addition, corporate users believe increasingly that packetized voice connections are more likely to survive disasters like September 11.

There is much less consensus on packetized voice for small and medium-sized businesses and residences. Some industry sources argue the following:

(1) Based on price advantage, small and medium-sized businesses will purchase packetized voice and Internet access over DSL as a less expensive alternative to purchasing multiple voice and data lines.

(2) Residential customers will buy a package of video, Internet access and packet voice services from a cable-television company. As a competitive response to cable, LECs will sell a comparable package.

The lack of consensus on small and medium-sized businesses and residential customers reflects the cost to migrate the embedded circuit-based network technology, the lack of agreement on QoS standards in public networks, and the absence of an RBOC champion for packetized voice.[16]

R&D Cutback Adverse in the Long Term

In order to conserve cash and survive, network-technology companies (in particular) and the industry (in general) decreased research and development (R&D) expenditures in 2001, will continue to do so in 2002, and may not increase such expenditures until 2003 or later. This will affect adversely the commercialization pipeline for new products and services during the 2003–2005 period (and potentially beyond). The cut back on R&D will affect disproportionately the future of small, independent technology companies that often combined seed money from established companies with other sources of capital.

Radical innovations that on average take 10 years between concept and market will be delayed or terminated prematurely. The full significance of this situation cannot be forecasted with precision, but its initial impact will occur in the last half of this decade.

The venture capital (VC) community's reluctance to fund the telecommunications-related technology sector reinforces the R&D shortfall effect. The VCs see (1) a declining demand for the technology and (2) no viable exit strategy once the capital is committed.[17]

Companies not involved currently in telecommunications directly (e.g., Microsoft) are *increasing* R&D expenditures. Such increases may provide competitive advantage in the event of entry.

To Regulate or Not To Regulate

The current situation could trigger a movement for more regulation, re-regulation, and/or legislative intervention. Such a movement will increase uncertainty and slow, if not stop, the industry's recovery process.

It is popular now for state and federal regulators to endorse "market-based solutions" in which regulators restrict their role to (1) ensuring a level playing field and (2) addressing areas where the market has failed. Given current market conditions, regulators have begun to identify "market failures" that potentially require some form of intervention. Examples of such alleged "failures" include the following

(1) A perceived "lack" of local competition as demonstrated by the financial failures of many CLECs
(2) The "low take rate" by small businesses and residences for "broadband" services
(3) "Underserved" rural and/or low income areas that do not have parity pricing and/or service availability with urban areas
(4) The apparent lack of spectrum for third-generation (3G) services
(5) The "excessively long" rollout of broadcast digital television

Regulators are not comfortable with the Schumpeterian concept that the current period is one of "creative destruction" in which the shakeout-consolidation process is natural and a prerequisite to moving the telecommunications industry from one growth cycle to another. In effect, the failure of many CLECs, structural consolidation, and "slow" adoption of "broadband" by consumers are interpreted by many regulators and legislators as the trigger for increased or re-regulation.[18]

The Global Perspective

Success in the United States constitutes a prerequisite for a foreign carrier to be considered "global" and therefore credible to home country and U.S. corporate customers. However, the potential for further foreign carrier investment in the United States remains highly unlikely in the 2002–2003 period.

Major foreign carriers remain unable to adjust to home-country competition and/or are crippled by the debt loads necessary to meet capital commitments (e.g., for 3G licenses and infrastructure construction). In addition, those that have tried U.S. entry strategies (e.g., BT) have made money on their stock purchases but failed to sustain presence in the U.S. market.

Experience to date shows that minority equity investments made in large U.S. carrier holding companies (e.g., DT in Sprint) do not meet the long-term business requirements of foreign carriers. A potentially more viable strategy consists of portfolio investments in market segments (e.g., wireless or information services) in which the investing carrier has experience and can achieve effective control (e.g., DT in VoiceStream).

"Global carriers" that focus on serving multinational companies can be expected, at a minimum, to provide telecommunications services between, among, and within the countries of "The Triad" (North America, Japan, and Europe). Carriers have attempted one of two global strategies: (1) "on-net" emphasizing end-to-end services using owned-and-operated facilities (e.g., WorldCom) or (2) a "partner" arrangement between or among two or more carriers (e.g., AT&T and BT in Concert).[19]

No carrier or carrier consortium can yet meet all the criteria for global service all of the time everywhere. However, some can meet many of the criteria most of the time somewhere. The critical differentiator is the ability of a carrier to provide price-competitive services *between, among, and within* the Triad countries. This may be accomplished by means of owned/leased facilities and/or the facilities of a partner and/or the facilities of a non-affiliated local carrier.[20]

Given the abundance of competition and the over-capacity of network facilities throughout the Triad, it may be that the real value-add is the identification and selection of local/regional providers and the aggregation of those providers into a network of networks to meet the requirements of specific customers. Consolidation of carriers in Europe and North America will facilitate the network of networks approach. Large business customers may move to self-provision rather than outsource under this model.

Managing for Success

The difference between corporate success and failure will come down to the ability of management to lead their enterprises through the downturn and position them for renewed growth. The telecommunications industry is critical to the U.S. economy and national productivity.[21] Therefore, after the current workout period, telecommunications will grow at least as fast as the overall economy if not faster.

The industry participants that contributed to *Strategic Trends* were asked to be descriptive not prescriptive. However, their comments often included clues and signals as to the strategies senior management would follow. The broad outline of these strategies is subsequently presented.

This is a period of triage in which the industry can be grouped into three tiers:

(1) *Tier I: Companies that are healthy fundamentally.* These should seize the opportunity to gain market-share and position themselves versus weakened competitors.

(2) *Tier II: Less-well-off companies that need to focus on survival and avoid risks.* For example, in the short run, a Tier II may focus regionally and forgo national expansion.

(3) *Tier III: The dead and the dying* that should sell out or fold up in such a way as to least injure their investors, employees, customers, and suppliers. To extend their lives at the cost of damaging the industry's reputation, wasting more capital, and providing an opportunity for re-regulation does the industry a disservice.

At the Tier I companies, management is leading along two dimensions: *the operational and the strategic*. In terms of *operational* matters, these companies emphasize back-to-basics. Core processes are being reviewed to ensure that the metrics are correct and standards are met (e.g., order-entry time, first-call customer problem resolution, mean time to repair [MTTR]). Personnel that cannot meet expectations are being terminated.

Tier I companies are moving to increase share while focusing on churn reduction and customer loyalty programs for the embedded customer base. A few are even seeking out customers that are experiencing short-turn recession-induced business effects to offer them extended payment terms and/or advantageous pricing in anticipation of a strong relationship in better times.

The Tier I companies are also renegotiating contracts and reviewing the number and performance of their sources of purchased goods and services. While not pushing suppliers to the wall, there are opportunities to drive costs out of the supply chain as well as lock down long-term supply relationships.

In the s*trategic* realm, this is the time to make *acquisitions* with asking prices low and going lower. Compared to 1999 and 2000, Tier I companies can acquire customers, assets, or new distribution channels at a fraction of the cost. However, acquisitions need to be planned carefully so as to truly add strategic value. The planning process needs to include post-acquisition implementation actions. Finally, given the over-capacity, poor management, and cash-flow problems in the industry, acquisitions must be selective and limited, or the process could easily destroy value and take management's attention away from the core business.

Tier I companies are leveraging their *internal competitive advantages* (e.g., scale economies to allow price reductions, access to capital to drive new technology deployment, increasing marketing to enhance brand and gain share when advertising prices are very low due to the recession). The key here is to use financial and commercial advantages to take advantage of temporary market aberrations and competitor distress.

After the Tier III companies and the weakest Tier II companies disappear, the industry will be better capitalized and more viable. Most importantly, after the shakeout, the telecommunications industry will be positioned to support a new era in wired and wireless communications. *The era will be all digital, increasingly video-focused, have intriguing new entrants (e.g., AOL and Microsoft), and be focused on anywhere, any-time connectivity supplied by a limited number of consolidated, well-financed major players.*

Notes

1. "Recovery" may be measured by such factors as increases in the rate of growth of revenue, operating income, cashflow, and capital expenditures (CAPEX) at the industry, segment, and/or company levels. Improved performance of industry stock indices will confirm a turnaround.
2. Local infrastructure in an acquiring RBOCs home region would be divested to meet antitrust requirements.
3. Market research shows a potential penetration level of up to 30 percent of urban households and 8 to 10 percent of businesses with 500 or fewer employees.
4. Preliminary comments by Comcast management at the announcement of the acquisition of AT&T Broadband indicate that the successor company will market and sell a video/Internet access/telephony service bundle aggressively.
5. While ILECs appear to have prevailed in voice-centric competition with the CLECs, they may have been victorious in the wrong war as (1) voice traffic migrates to wireless and (2) wireline competition becomes data-centric.
6. "Predominately facilities-based" includes the operating model in which a CLEC leases local-exchange carrier (LEC) unbundled loops on a long-term or short-term basis, depending on the economics of the transaction.
7. Distressed CLEC assets can be purchased for 20¢ or less for each dollar on the CLECs balance sheet.
8. Federal Communications Commission (FCC) statistics show 5.2 million "high-speed" connections (i.e., 200 kbps or better) at year-end 2000, of which 64 percent were cable modems, 31 percent ADSL, and 5 percent other. Based on these numbers, only 4 to 5 percent of U.S. households had high-speed connections at the end of 2000.
9. The official position of the FCC is that availability, not take rate, is the key measure.
10. The range of potential government interventions includes (1) tax subsidies for high-speed access providers and/or homeowners, (2) direct subsidies similar to the Universal Service Fund, (3) removal of restraints such as problems with digital rights management (DRM) and rights-of-way, and (4) government responsibility for construction and operation of networks especially in rural areas.
11. The September 11 attacks and the current war on terrorism make DOD spectrum untouchable in the near term. Potentially, some spectrum may become available through one or more of the following: (1) early migration by broadcasters out of channels 60–69, (2) agreement on the use of NextWave's spectrum, and/or (3) agreement on mobile use of multichannel multipoint distribution system (MMDS) spectrum.
12. The "surge era" refers to the excess of aggregate capital spending on network construction and retrofit that occurred during the 1996–2000 period over the baseline 1991–1995 period in the United States.
13. For example, carriers with embedded circuit -switched networks are disinclined to procure unproven softswitch technology when Class-5 switches have sold recently at less than twenty cents on the dollar.
14. As discussed previously, for small businesses and consumers, the growth rate for broadband access has begun to decline. That limits the pace of network investment in upgrading the local loop.
15. Such home networks may be wired or wireless, hopefully using one standard agreed to by participating companies so as to not confuse consumers with incompatible subsystems that cannot be interconnected. Home network applications include (1) security, (2) entertainment, (3) home automation, (4) PC networking, and (5) intra-home communications with outside-home connectivity.
16. In November 2001, the local Telecommunications Division of Sprint announced a contract with Nortel Networks to transform its network to a packet voice network. The Division has eight million customers across 18 states. This is the first move by an LEC to transform its embedded network to packet-based technology.
17. An initial public offering (IPO) or acquisition constitutes an acceptable exit strategy. Neither option appears realistic currently.
18. In the near term, carriers that trust the market more than the regulatory/legislative process will put relatively more personnel and resources into regulatory reform/neutralization than was the case in the 1996–2000 period.
19. Neither strategy has proven to be superior. The "on-net" approach is very capital intensive, while "partnering" requires constant management and refreshment of inter-company alliances.
20. To be price competitive when requiring facilities from a non-affiliated local carrier, a carrier must negotiate substantial inter-carrier price discounts (i.e., a wholesale rate based usually on minimum volume commitments).
21. In a speech to investment analysts (December 4, 2001), Sprint Chairman and Chief Executive Officer William Esrey commented that telecommunications "constitutes two percent of the $10 trillion U.S. economy" and that he expected telecommunications revenues to increase at a seven percent compound annual growth rate between 2001 and 2005.

Broadband—Hype and Reality

Balaji Kumar
Director, Advanced Network Services
EngineX Networks

Abstract

This paper provides a unique perspective on the state of the broadband industry and its service deployment. It gives a snapshot of the broadband promise made by industry participants and assesses the current state of that promise. In addition, this paper discusses some strategies on what it will take for companies to achieve the broadband promise. It also provides some insights, which should help all of the players in the food chain to benefit in terms of profiting from the broadband service deployment.

The Broadband Promise

It all started out with a simple "broadband" promise by the communications industry (service providers, equipment vendors, and others) to end users. The promise was to provide unlimited on-demand bandwidth, a choice of providers, cost-effective services (voice, data, and video) with high reliability and security at real broadband speeds. In the process, industry experts forecasted that broadband with converged services would generate billions of dollars in revenue. According to CIMI Corporation, worldwide network-based service providers are expected to gain incremental revenue streams totaling more than $600 billion during the next ten years. This revenue could translate into huge profits for the providers, while fostering the creation of an array of services benefiting end users. In this environment, end users will have the flexibility in communications, information access, and other services. This will definitely change the way people live and communicate.

To enable this promise, the industry has taken actions in multiple directions—technology, regulatory, and business financing—to accelerate the goal of achieving the broadband promise.

The Broadband Reality

However, the seemingly simple broadband promise has become a nightmare for the communications industry—service providers, equipment vendors, investors, and users alike. There is no question that there is a huge demand for broadband. In fact, according to Cahners In-stat (see *Figure 1*), broadband access (a primary indication of broadband demand) is expected to penetrate 32 million households (30 percent of U.S. households) by 2004.

Many believe that broadband has arrived because there is a huge demand. However, the underlying reality paints a very different scenario.

The whole industry is collapsing due to the cumulative effect of competitors failing to deliver the technology or service as committed and the inability to meet regulatory mandates and access requirements, which are together compounded by the Internet's current state-of-technology. This has stunted the growth of business models that profit from broadband delivery. The problem that hinders the growth of broadband falls into five categories:

1. Business Model
2. Technology
3. People
4. Finance
5. Services

Business Model
A sound business model is the basis for any business—there is no exception granted to broadband. This means a model with which a business can sell products or services that customers are willing to pay for and, in the process, generate revenue and be profitable over time. The service-provider industry, which is pursuing the broadband opportunity, consists of the incumbent local-exchange carriers (ILECs)—carriers that spun-off from AT&T since divesture in 1984—and new entrants leveraging existing and new technology for more efficient service delivery. The business mindset and the requirement for these segments are completely different. The incumbents think long term for return on investment (ROI) because they are not in a hurry to fill customer needs or demand. Their mindset is that "the user will buy it when we build it." New entrants, on the other hand, have a short ROI mindset, usually driven by the time frame and expectations of their investors. The new entrants that raise capital from the public have to show that they are different from others and use technology as an excuse to differentiate their value proposition—at least they did so when market valuations were reaching new heights daily. Examples of such companies are Teligent, Winstar, Covad, Rhythms, and Northpoint. Because these players used technology as the pitch, it forced the equipment vendors to pitch product and services that providers were demanding. The reality was that these technologies—which were the basis for the companies' business models and very reason for being—were either premature or incomplete, or did not fit into the

FIGURE 1

Broadband Subscription

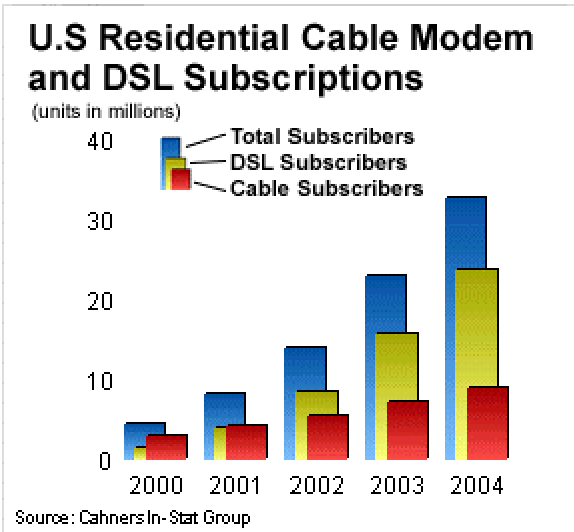

providers' service offerings to their customers. In addition, the costs to develop these technologies far exceeded the revenue they could generate for the providers.

Thus, the industry began with a model that was built on technology that had not been validated in terms of user needs or applicability of the technology to the pressing business needs of service providers. But, the industry's hype and buzz about broadband in the heady, gold-rush era (pre-bust) meant that the fundamentals with respect to building a business were ignored.

Technology

Technology was actually the driving force behind the broadband promise. It brought down the barriers with respect to cost, speed, and time to market and allowed more competition to the industry. With the prospect that the communications industry could radically change everything about broadband delivery, many companies joined the cause.

However, too much of anything is always a sign of disaster. There is more technology available in the communication industry than anyone can handle. There is asymmetric digital subscriber line (ADSL), high–bit rate DSL (HDSL), very-high–data rate DSL (VDSL), G.shdsl, local multipoint distribution system (LMDS), multichannel multipoint distribution system (MMDS), wireless application protocol (WAP), asynchronous transfer mode (ATM), and Internet protocol (IP), to name a few. In a majority of cases, these technologies compete against one another without any unique value from the end user's point of view.

In addition, those with financial incentives in the industry—bankers, media, and instant experts—hyped the technology beyond its capabilities. Anyone who visited a trade show and could spell one of the alphabet soups was considered an expert. When this kind of expert controls the capital or influences people who control the capital, the chain reaction becomes uncontrollable. One technology would be anointed, and that year all the investment would follow into the sector that develops that technology or to the provider that provides the services using that technology. The most recent trend in 2000 was the broadband wireless MMDS. This is the last of the broadband technology, which has its limitations as any other previous broadband technologies, such as DSL, cable, and LMDS. It was expected to solve all of the previous problems and achieve the broadband promise. What the innocent fails to understand is that wireless technology is one piece of the puzzle in delivering the service to the user, and it is in no way a "be-all, end-all" solution.

People (Management)

It does not matter how credible the business model or how high the market demand is for the service/product or the technology. It is the experience and credibility of the people behind it that makes the difference.

The industry had promised users what they could deliver based on the hype, and the users did not have the skill set or understanding to differentiate reality from hype. The challenge now faced by management is a combination of traditional thinking with a new business model (aggressive unproven model) or developing a new technology, or simply the raft of start-up environment challenges. In addition, difficulty in hiring regulatory and financial people with limited knowledge of technology (or the telecom industry) further complicates the situation. The bottom line is that many management teams are not capable of understanding the end-to-end issues pertinent to successfully executing something so complex, especially when there is no precedent.

Today, the road is littered with numerous examples of what not to do. What is less prevalent are examples of companies with a track record for doing it right. There is no best practice to follow. In addition, human nature is to learn from mistakes and not to avoid problems beforehand by logic.

Therefore, the challenges with respect to people capable of delivering the broadband promise are far from resolved.

Finance

Because the broadband technology is new to the industry (less than 15 years) and the communications business is a capital-intensive one, the new players require huge amounts of capital to build either a product or a network to provide uniform service to end users. It also requires the time to develop the product/service and a competent team to deliver the vision, as well as patience to monitor the progress using valid benchmarks to measure progress.

The financial community does not understand the broadband environment and expects the project to be completed in half the time with half of the cost. The folly of this unrealistic set of expectations is painfully clear in the case of the demise of the telecom sector. Broadband Office, Covad, Rhythms, Northpoint, Teligent, and numerous satellite-based companies raised millions, if not billions, of dollars with business models that depended on unproven technology. In addition to this, the industry was developing products without any understanding of customer needs or addressing the issue of affordability. Thus, the industry spent millions without having any tangible return to investors. The inevitable bubble bust surprised the experts tracking the sector—but they really should not have been surprised if they had applied critical thinking in their company/industry analysis.

Services
One of the biggest promises of broadband was to deliver a wide variety of services—such as voice, data, and video—at high-speeds and at a fraction of the cost compared to the existing solution.

The reality is that the services were provided using unbundled network elements (UNEs) whose total cost of providing the service exceeded the price charged for the services. The high cost was incurred due to the services being delivered using a "stove-pipe model" (i.e., a vertical integrated system and process to cater to a specific service offering, as in the case of today's telephone service). This forced the providers to provide either a limited set of services or unmanaged services for which the users were not willing to pay any premiums, which would have helped the providers cover the cost of the infrastructure build-out.

Moreover, competitive pressure and regulatory complexities have made the service offering an expensive solution such that only companies with high capital and a larger customer base could survive and keep providing service at low or no margin.

Figure 2 shows a typical broadband environment for a service offering. Here, the majority of the broadband (services such as voice, on-demand high-speed access to information) is delivered using a variety of access technologies such as DSL and T1 with dedicated links via ATM/IP to the appropriate backbone provider with no bandwidth guarantee or security or quality of service (QoS).

The customer who subscribes to network services gets Layer-2/Layer-3 transport service from point A to point B. The customer pays the provider a flat fee for the use of this transport to access anything. (The customer pays irrespective of the use of the transport.)

Thus, the provider gets a flat fee irrespective of the bandwidth used by the customer. The provider is in no position to leverage the unused capacity to generate additional revenue. This is because the provider does not control the bandwidth use or sell service-level agreements (SLAs). The providers do provide SLA–based services today by offering dedicated private lines or by charging based on the content, service, or application used or accessed on a demand basis.

Because of the various network configurations and technologies used, providers are unable to provide advanced services. One of the reasons for the voice network being very successful is that the provider has complete control of the services delivered over the network and is thus able to charge or give away services from time to time based on

FIGURE 2

Today's Broadband Environment

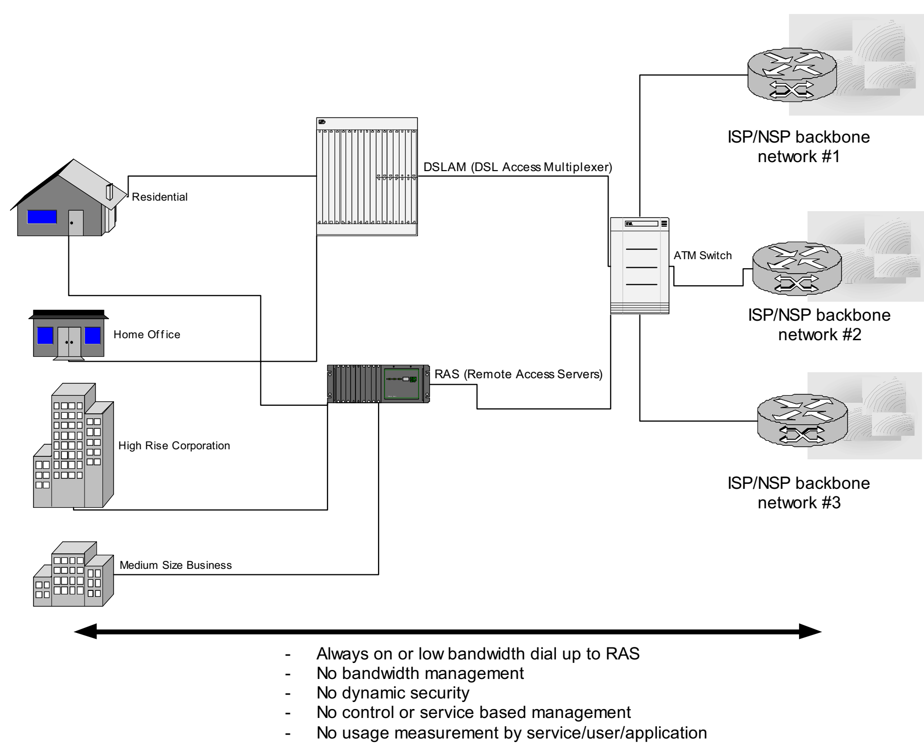

demand and usage. This enables value-added services, such as caller ID, fax, and dial-up options, to be charged for at a reasonable cost while being profitable to the provider.

Target Environment for Fulfilling the Broadband Promise

The previous section discussed the current reality of the broadband environment, where there are numerous challenges in achieving the promise envisioned. Even though it is not an impossible task, companies are still a long way from achieving it.

In this section, we will discuss the same categories mentioned earlier with the appropriate strategies that can be pursued in order to fulfill the broadband promise.

Business Model

Irrespective of the type of business, there are certain expectations by the investors with respect to ROI. The private investors, who invest in the new players, have a very high expectation with respect to short-duration ROI. In the quest for capital, new players responded by building business models that met the investors' unrealistic expectations. This, in recent years, has ended in disaster. A typical rule of thumb for a broadband service-provider business model focused on convergent services is payback within five years with at least 5 percent penetration or take-rate for a given serving area, assuming that there is sufficient automation from day one to keep the cost down. If the business model deviates from this rule of thumb, it would be wise to validate the assumptions and capability of the team in executing the model. For instance, service providers that plan to provide the services need to validate these services with the currently used technology (mature technology is low risk), alternatives for providing the same services, take-rate assumptions, coverage, customer willingness to pay for the service, etc. It is necessary for these parameters to be validated with sufficient scrutiny. This litmus test should indicate that a model is built to withstand even the worst times.

Technology

One of the primary challenges faced by the industry is building a business model around a technology. For instance, Teligent and Winstar were classified as broadband wireless, whereas Covad, Northpoint, Rhythms were classified as DSL players, even though they were targeting the same business and service offering. The obvious reason was that the investment community was dazzled by the technology promises without any basic business sense.

The basic recommendation is that the business should not be driven by technology. Instead, the business should be market-driven based on customer needs and requirements and not based on the promises made by technology capabilities. Thus, the service offering should meet customer needs such as price, QoS, better customer services, and wide range of service set. Technology should be an enabler for these service offerings.

In addition, in the case of unproven technology in the field, the first rule of thumb is to avoid it. If one is willing to take the risk, then the technology should be validated in various controlled environments before being rolled out to the customer. Part of the validation should be an end-to-end fit of the technology as a solution in the provider environment along with other existing systems. This enables the provider to see the real cost savings and benefit for the company and its customers. One such general focus should be automation, which reduces the manual intervention as much as possible.

Nevertheless, this is more easily said than done. Due to the pressure from investors and the promises made by the management team, new players in the industry have tried to short cut these efforts, with predictably dismal results. Only the large incumbent players with built-in support could sustain such inefficiencies.

People

This is the most complex area to address and is the one that makes or breaks a company's ability to execute on its vision to provide broadband successfully. It is management's challenge to build a team with the right abilities and attitude to execute the vision. The basic objective of any team should be defined clearly so that everyone on board has realistic business goals that can be achieved or executed. If the business goals are based on promises that are not realistic, then the team will be demoralized or try to cut corners to meet the objective.

It is necessary to have a team that is level headed and experienced, with a proven track record for thinking out of the box, using some logic, and setting realistic goals and expectations.

Finance

One fundamental rule to impress upon investors interested in the broadband arena is that there is no quick return on their investment. Long-term commitment is needed for emerging technology. The focus should be on solutions, and the role of technology should be as an enabler with respect to the automation of end-to-end service delivery with respect to systems and processes. Equally important, the investor should invest in a credible team to execute the vision.

Services

As mentioned earlier, the services should be driven by the market demand and to meet customer needs. Some of the basic customer needs are low cost, ability to track and measure services; and multiple services (voice, data, video) using a common infrastructure.

Figure 3 shows a target broadband network environment that meets the basic customer requirement with respect to the flexible, cost-effective delivery of services such as voice, data, and video. In this environment, the customer negotiates dynamically with the provider regarding requirements, such as services (voice, data or video); expected bandwidth; and expected QoS. In this architecture, the provider will connect to a service switch in the network and negotiate the network parameters, and provide the necessary access to information or services that the customer needs. This architecture also enables a provider to track and measure the services being offered. This will enable the provider to charge or bill appropriately for the various services delivered to the end user. This architecture is very similar to the voice (narrowband) network but much more complex.

Summary

This paper has examined the state of the broadband industry and its service deployment compared to the promise of broadband. The broadband promise was to provide unlimited on-demand bandwidth, choice of providers, and cost-effective services (voice, data, and video) with high reliability at real broadband speeds. In the process, the industry would generate billions of dollars of revenue with converged services.

What started as a simple promise has become a nightmare for the communication industry—service providers, equipment vendors, investors, and end users. The challenges range from developing appropriate business models, building reasonable assumptions about the nature of the technology solution, the need to hire experienced and talented people, setting reasonable and measurable goals that can be articulated to the investors funding the companies, and, last but not the least, the services offered. Failure to meet these challenges has brought broadband to its current reality.

In general, there are no best practices to follow at present, but there are some high-level litmus tests that can be done vis à vis the business model, technology used, staffing, investor objectives, and services offered. Some are generic; others are specific to broadband environment. Any amount of progress in these areas will help the industry to progress toward its promise.

FIGURE 3

Target Environment for Fulfilling the Broadband Promise

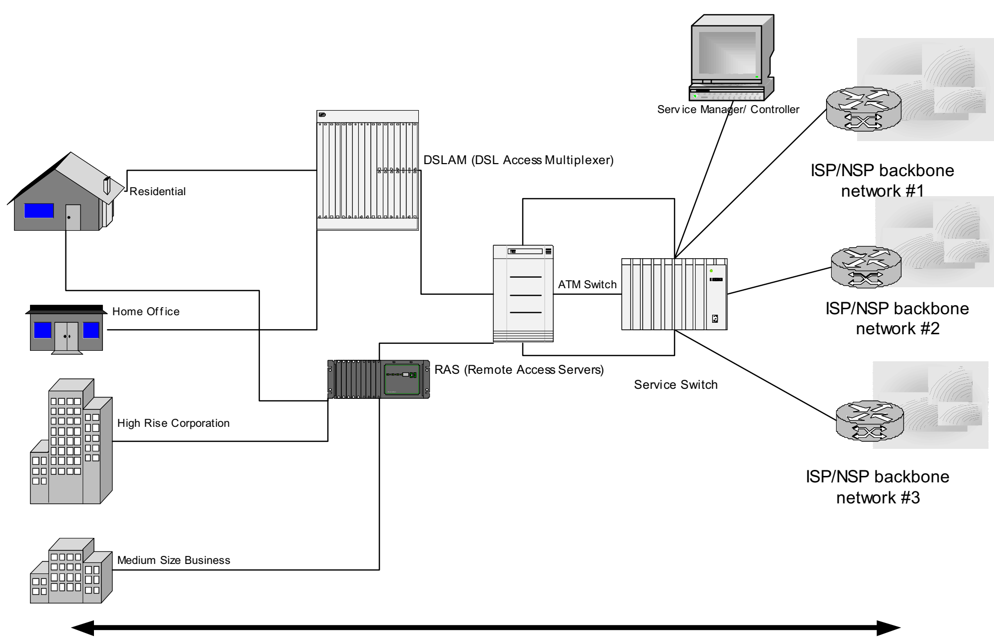

- On Demand bandwidth either broadband or dial up to RAS
- Bandwidth management by customer/ service/ application
- Dynamic security / Firewall / VPN
- Control or service based measurement
- Usage measurement by service/ user/ application

The Business Case for Moving Voice and Data Concentration to the Edge

John Maddison
Director, Product Marketing
VINA Technologies

Carriers have long known how to capture revenue by deploying voice and data services in metro access networks. But to do so profitably remains a formidable and perhaps even growing challenge. As recent revelations of red ink and shrinking margins make painfully clear, providers have especially struggled to make money when they extended bundled services to small and medium-sized businesses (SMBs)—a vast and underserved market. Earnings have often eluded the sector's participants because they lacked the technology to optimize their existing infrastructures and to connect customers affordably.

To maximize the productivity of their networks and to control the cost of adding new accounts, years ago carriers adopted the tactic of concentrating their voice and data. The first type of concentration relied on digital loop carriers (DLCs), while the second took advantage of digital access and cross-connect systems (DACSs). Most of the companies aimed to concentrate their voice by 4:1 and their data by 3:1. Both ratios reflected the statistically proven theory that only a fourth of the voice subscribers and just a third of the data subscribers need access to a network at the same time. Meanwhile, a different kind of concentration was going on in central offices (COs), where carriers were increasingly clustering their DLCs and DACSs. With the instruments of their voice and data concentration confined to a relatively few large sites, the providers hoped to shrink their spending for capital equipment and thus aid their bottom lines.

But a funny thing happened on the way to profitability. The emphasis on centralization never produced its expected savings. In fact, the cost of connecting customers to networks steadily rose. To make matters worse, the increase coincided with an explosion in demand for bandwidth and with a resulting growth in required capacity per connection.

What went wrong? For one thing, the decision to make COs the focus for concentrating services unwittingly compounded carriers' requirements for T3 or larger trunks, each often leasing for at least four figures per month. Then, too, the attempted solution never adequately accounted for the proliferation of the DACSs and DLCs that proved necessary to support the providers' CO–based Class-5 switches and high-capacity routers. Moreover, despite concerted efforts to centralize the concentration of voice and data, some of the DACSs and DLCs wound up in co-locations or local serving offices anyways. To the extent that they did so, the devices hurt incumbent vendors by aggravating their shortage of unused racks and hampered competitive providers by inflating both their capital and recurring spending, including their fees for leased space. From both an operational and a capital perspective, therefore, the original paradigm for concentrating broadband communications backfired. Since then, the cost of connecting customers to metro access networks has continually climbed without showing any sign of ending.

Growing financial pressures are increasingly prompting carriers to consider alternative approaches for bringing new end users on-line. One proposed solution that shows particular promise consists partly of transplanting time division multiplexing (TDM) DACSs, GR303 DLCs, and frame-relay switches from COs to co-locations and thus placing the products as closely as possible to customer premises. The proposal further involves combining the devices under one cover or even on the same card to form a new kind of system that for the first time concentrates voice and data simultaneously. Some industry analysts have recently taken to referring to the invention as an integrated multiservice access platform (IMAP).

In different respects, both parts of the contemplated innovation hold the potential to help carriers substantially decrease their spending and improve their earnings. Transferring DACSs, DLCs, and frame-relay switches—and therefore the concentration of voice and data—to the edge of metro access networks would reduce operational expenses by optimizing backhaul capacity and lowering demands for T3 or bigger trunks. At the same time, integrating all three hardware modules to produce an IMAP would slash capital spending by replacing them with one multifunctional unit that precludes their purchase as discrete footprints. The integration of formerly separate boxes would also bring the additional benefit of saving room in racks and minimizing monthly bills to lease space in co-locations.

In short, the combination of extensive integration and decentralized concentration promises dramatically diminished capital and operating expenditures and greatly expanded opportunities for profitably adding customers to carriers' networks. Confirmation of the breakthrough comes from early adopters of newly introduced IMAPs, as the rest of this white paper shows. The following three brief case histories present a compelling business case for repositioning the concentration of telecommunications traffic from COs to a multifunctional platform at the network's edge. As the report additionally illustrates, the emerging model's overwhelming benefits apply to competitive and incumbent vendors alike and to both voice and data services.

Example #1: A Newcomer Saves Big by Concentrating Its Voice

At one large competitive carrier, IMAPs cut capital and operational spending over two years by more than $31 million.

Delivering voice and data services in 32 geographic markets, the carrier serves SMBs that average 10 lines for phones and six for connection to the Internet. Integrated access devices (IADs) in each of the end customers' premises converge the parallel telecommunications signals onto a single T1 link that leads to one of the company's 750 co-locations. Before acquiring its IMAP, the carrier equipped every co-location with three M13 multiplexers, each accepting up to 28 incoming T1s and aggregating them to a T3 trunk. Otherwise, the company centralized its telecommunications gear as much as possible in the hope of squeezing costs through economies of scale. Toward that end, each of the provider's COs housed a DACS that steered the T3's voice traffic over 28 T1s to a coresident Class-5 switch and the data across 18 T1s to a coexisting router. The Class-5 switch and the router then respectively forwarded their transmissions to the public switched telephone network (PSTN) and to a frame-relay network.

Under the pre–IMAP regime, limitations in the M13s' capacity forced their owner to install a new multiplexer after every 28th customer. The resulting proliferation of M13s in the company's co-locations inflated expenditures for capital hardware. Worse still, each additional multiplexer also raised the carrier's recurring costs by necessitating another T3 backhaul connection, leasing for $2,500 per month.

To restrain its expenses, the company replaced its existing DACSs and M13s with IMAPs in all its co-locations and configured each platform so that 68 of its 96 T1 ports serve customers. The remaining 28 ports act as backhaul connections. Because such a configuration allows for nearly 2.5 times more customer-facing broadband links than an M13, the carrier can now do away with all three of the multiplexers that previously operated in each of its co-locations. At $3,000 each, the three eliminated M13s mean a reduction of $9,000 per co-location in outlays of capital equipment and a saving of $6.8 million across the company's entire network of 750 sites.

By relying on the IMAP's ability to concentrate voice signals, the carrier similarly lowers its requirements for backhaul trunks from three T3s per co-location to just one. The removal of two links, each leasing for $2,500 monthly, at all 750 co-locations whittles the provider's backhaul costs by nearly $1.9 million per month or $45 million over two years.

For every eliminated T3, the company also conserves a port in its CO–based DACSs. So two retired backhaul trunks at each of the 750 co-locations translates to 1,500 saved DACS ports at $1,000 per unit, and thus further decreases the carrier's capital spending by $1.5 million.

An additional six-figure drop in capital expenditures results from the IMAP's ability to optimize the usage of ports in the

FIGURE 1

An IMAP Simplifies a Start-Up Carrier's Network and Cuts Its Costs

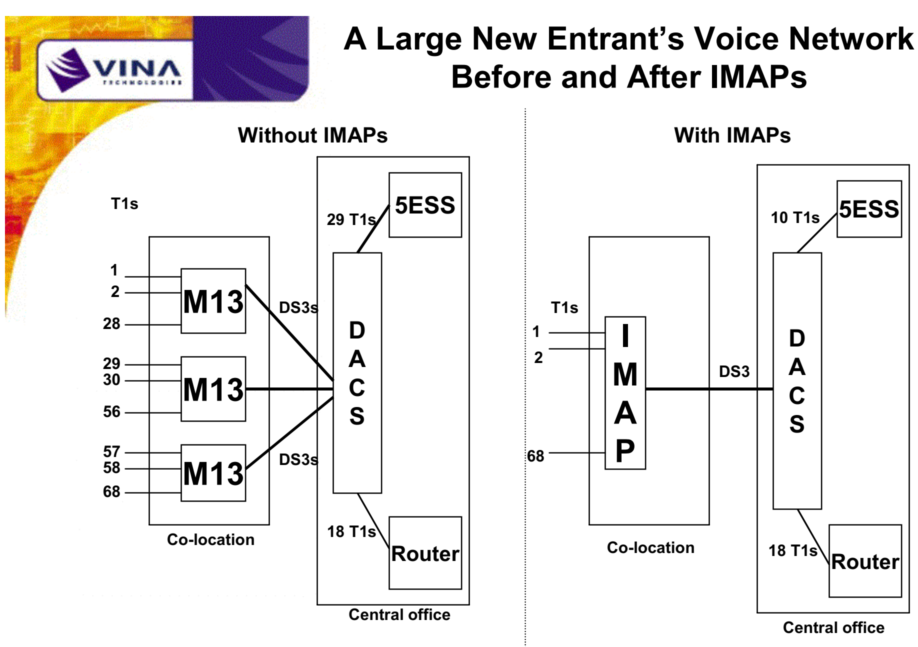

service provider's Class-5 switches. Using the GR303 protocol, the platform concentrates the carrier's voice traffic by threefold and thereby enables it to trim the number of T1s between the switches and their supporting DACS from 28 to 10. The conservation of 18 T1 connections, equivalent to 440 digital signal (DS)–0 ports in Class-5 switches at $100 per port, saves the provider an extra $44,000 at each of its COs and $1.4 million for all 32 of them.

The decline in outlays for switches brings the company's gross drop in payments for capital and recurring expense items to $52.4 million over two years. Included in the total are savings of $1.5 million for DACS ports, $45 million for leased backhaul lines, and $4.5 million for M13s. After accounting for the IMAPs' aggregate cost of roughly $23.3 million, the provider observed a two-year net decline in spending of $31.4 million. By allowing the carrier to slash its installed T3 trunks, Class-5 switches, and DACS ports, the platforms also triple the useful life of the company's voice network.

Example #2: Carrier Saves in Converged Voice and Data Networks

Roughly 20 percent of the services that a small northeastern competitive carrier delivers to 208 SMBs consist of voice alone; the remaining 80 percent combine telephony with data. The customers average 10 lines for voice plus six for access to the Internet and converge their traffic with the help of an IAD onto a single T1 connection.

Either way, before the provider switched to IMAPs, each of the resulting 208 T1s went to one of the provider's 30 co-locations, where a DACS accepted the traffic and directed its voice component to a nearby DLC and its data to a frame-relay switch. After being concentrated four to one by the DLC, the voice transmissions proceeded to as many as three co-resident M13 multiplexers that aggregated the incoming T1s to T3s. The data, meanwhile, passed through the frame-relay switch without benefit of concentration and, like the voice, continued to the M13s. From there, both kinds of signals emerged from the co-locations and traveled over three backhaul trunks to a CO, where a Class-5 switch routed the voice to a PSTN. At the same time, the data moved first through the CO's core router, then to a frame-relay network.

Because of its past absence of a cost-effective means for concentrating data, the carrier found itself forced to maintain more extensive backhaul connections than it deemed optimal. In fact, to support its 208 incoming links from SMBs, the provider needed three backhaul T3s, collectively equivalent to 22 T1s for the concentrated voice and 42 for data, at each of its co-locations. Leasing monthly for a unit charge of $1,800, the links cost the carrier more than $1.94 million per year in recurring fees across its entire network.

Another source of unacceptably high operational expenses involved the company's monthly payments for leased space in its co-locations. To accommodate its broad assortment of single-purpose devices—such as DACSs, DLCs, and data switches—the carrier had to pay for two seven-foot racks per co-locations. Each rack leased for $800 per month to produce an annual network-wide bill of $570,600 for all 30 sites.

For most carriers, heavy spending for leased space in co-locations usually implies large outlays for capital equipment, and the subject of this case history proved to be no exception to the rule. To support its existing base of customers, the provider needed to equip each of its co-locations with four DACSs and DLCs and with enough ports in frame-relay switches for 998 DS–0s. The DACSs and DLCs averaged $36,720 in price, while the switches went for $23 per DS–0, equivalent to nearly $23,000 for every co-location. Moreover, the carrier required a minimum of three M13s at $3,500 per unit in each co-location to serve its end users' needs. The necessary number of multiplexers reflected their limited capacity of just 28 T1 connections per unit.

Concerned that its legacy infrastructure and associated overhead were fast spiraling out of control, the carrier started in early 2001 to consider a fundamentally new strategy for deploying converged services in metro access networks. The rethinking ultimately led the company to reconnect its customers' 208 in-bound T1s to IMAPs in each of its co-locations. Installation of the platforms signaled at least

FIGURE 2

Summary of an IMAP's Benefits for a Start-Up Carrier's Voice Network

A New Entrant Cuts Its Capital and Operational Costs by Using IMAPs for Voice

Asset	Saved Units	Unit Price	Sites	$ Savings
M13s	3	$3,000	750 co-los	$6,800,000
DACS ports	2	$1,000	750 co-los	$1,500,000
DS–0 ports	440	$100	32 markets	$1,408,000
All cap. equip.				$9,708,000
T3 trunks	1	$2,500/month x 24	750 co-los	$45,000,000
Gross CAPEX/OPEX				$54,708,000
IMAPs				–$23,250,000
Net CAPEX/OPEX				$31,458,000

two major advances in the provider's network infrastructure. First, the IMAPs' integrated DACSs and DLCs made their discrete counterparts in the company's existing leased space functionally redundant, thereby prompting their replacement. Second, the IMAPs corrected a previous glaring weakness in the carrier's technology—the omission of data concentration. Thanks to the platform's built-in frame-relay switches, the provider gained the newfound ability to concentrate its data by a factor of three.

In the wake of recent improvements, the carrier's network-bound transmissions now take the following revised path: After accepting traffic from SMBs, each of the IMAP's service modules processes voice and data respectively in its integrated GR303 DLC and frame-relay switch. While GR303 replicates the fourfold voice concentration that the provider practiced before, the switch concentrates its data 3:1. Both components of the solution then forward their respective loads to T3s that connect the co-locations to a CO.

By expanding the carrier's existing concentration of voice to include data, the IMAP shaves from 64 to 40 each co-location's required number of leased backhaul T1s for both services. This expansion also reduces the need for T3s from three to two. Of the 40 current T1s, only 16 correspond to data, a form of traffic that formerly demanded 42. The retirement of one T3 at each of the provider's co-locations translates to a monthly slice in recurring expenses of $1,800 per site and to a network-wide decline of $640,800 each year. Threefold concentration of data also lowers the company's requirements for ports in data switches from 998 DS–0s to 371 per co-location. At $23 per DS–0 in 30 facilities, the difference yields a total additional savings of more than $430,000.

Even more dramatic than the IMAP's impact on operational spending is its lowering of the carrier's capital outlays and the resulting conservation of the company's precious cash. Because of its multifunctionality, the platform frees the carrier from having to pay a unit price of $36,720 for the four DACSs and DLCs that reside in each of its 30 co-locations. The ensuing decrease in payments for hardware, therefore, totals $146,880 per co-location and more than $4.4 million across the company's full network.

A sharp cutback in the use of data switch ports plus the wholesale retirement of discrete DACSs, DLCs, and multiplexers diminishes the need for leased space in co-locations. In fact, since taking delivery of its IMAPs, the carrier has halved its monthly fees for rented real estate from $1,600 to $800 as its consumption of seven-foot racks has shrunk from two per site to one.

Spread over a whole year, the projected drop in expenditures will total $936,000 for operational items and $5,154,030 for capital hardware. In comparison to previous levels of spending, the two sums respectively save the carrier 61 percent and 31 percent.

Example #3: A Carrier Shrinks Its Capital Cost without Concentration

As the two preceding case histories attest, much of an IMAP's value rests on the product's potential to concentrate voice and data or just voice alone. But even carriers that shun concentration can still benefit mightily from the platform's multifunctionality and the resulting ability to dispense with duplicate hardware and minimize capital procurements. A striking case in point involves a large

FIGURE 3

How an IMAP Transformed a Voice and Data Network

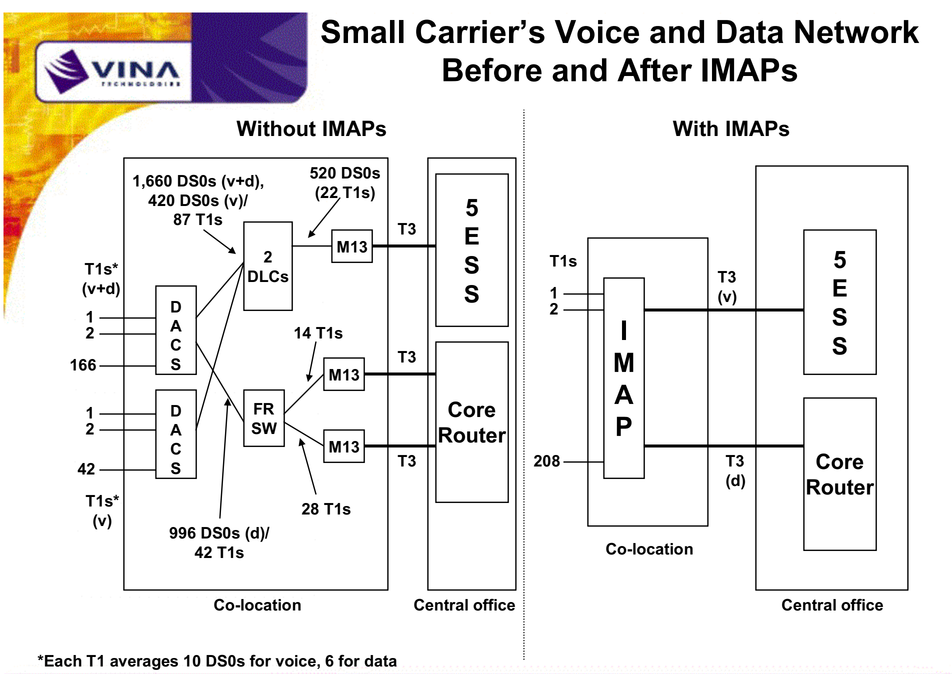

FIGURE 4

Summary of an IMAP's Value Proposition in a Voice and Data Network

A Small Carrier Cuts Its CAPEX and POEX by Using IMAPs for Voice and Data				
Asset	**Saved Units**	**Unit Price**	**Co-locations**	**$ Savings**
M13s	3	$3,500	30	$315,000
DACSs/DLCs	4	$36,720	30	$4,406,400
FR switch ports	627 DS–0s	$23/DS–0	30	$432,630
All cap. equip.				$5,154,030
T3 trunks	1	$1,800/month x 12	30	$648,000
Co-lo racks	1	$800/month x 12	30	$288,000
All OPEX				$936,000
Gross CAPEX/OPEX				$6,090,030
IMAPs				–$3,300,000
Net CAPEX/OPEX				$2,790,030

midwestern provider that differs from both of the earlier subjects in qualifying as an incumbent vendor and in declining to concentrate its traffic. Like its two competitive counterparts, the incumbent had grown increasingly disgruntled with the costliness of its existing model for deploying services and had recognized the need for fundamental change.

In one state with 25 metro markets, each typically comprising 288 SMBs, the carrier delivers a combination of services that demand an average of 12 lines per customer, eight for voice and four for data. An IAD in every served site integrates parallel lines onto one T1 connection that ends in the provider's local serving office. Because every end user corresponds to a different broadband link, the number of arriving T1s totals 288.

Prior to the carrier's purchase of IMAPs, all of the incoming T1 lines connected through a channelized service to 336 channel banks in each of the 25 metro locations. The number of installed channel banks reflected the need to maintain one for each SMB and another 48 to support the average requirement of four data-oriented DS–0s per customer. After reaching the channel banks, the voice and data signals diverged—the former heading to a Class-5 switch for transfer to the PSTN and the latter bound for a DACS en route to the carrier's data network. Half of the fully loaded DACS's

FIGURE 5

A Large Incumbent Carrier Adds an IMAP without Concentrating Its Traffic

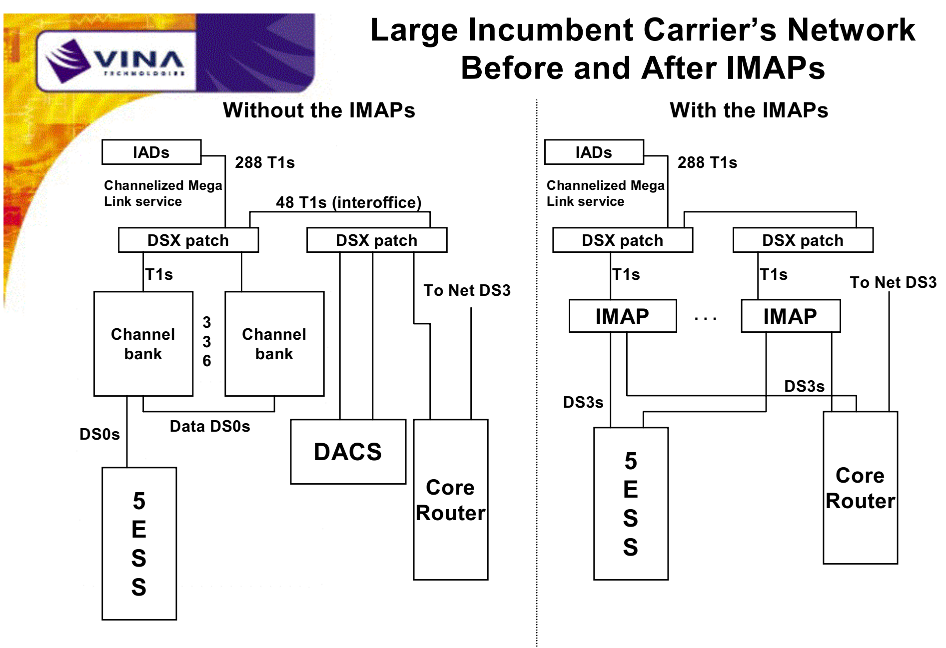

FIGURE 6

Summary of an IMAP's Value Proposition for a Large Incumbent Carrier

An Incumbent Slashes Its Capital Spending by Using IMAP But No Concentration				
Asset	**Saved Units**	**Unit Price**	**Sites**	**$ Savings**
Channel bank	336 per site	$2,000	25	$16,800,000
DACS	1 per site	$20,000	25	$500,000
CAPEX cost				$17,300,000
MBXs			25	–$5,570,750
Net cap. effect				$11,729,250

96 ports were allocated to outgoing T1s from end users. The remaining 48 accommodated T1s that led to the company's data network through an intervening frame-relay switch.

Even under the most generous terms, however, the company faced the intimidating prospect of footing a bill of $2,000 for each of the 336 channel banks in every market. Total amount due per site: $672,000, excluding the additional budgetary bite of $20,000 for the provider's data-handling DACS. In short, the fiscal drain from its capital investments struck the provider as unjustifiably severe.

Convinced that overturning its networking status quo would reap huge dividends, the incumbent removed its installed DACSs and channel banks in favor of a few IMAPs that functionally emulate the devices that they replace. As a result of the change in infrastructure, converged T1 traffic that once traveled to the channel banks now goes instead to IMAPs over the same channelized link. From the new platforms, the voice and data then move as before to a Class-5 switch or to a frame-relay switch without the slightest aid of concentration.

Since the IMAPs' arrival, the carrier's outlays for capital equipment have plummeted almost 68 percent. Most of the plunge results from the platforms' threefold lower cost than the supplanted channel banks: In each metro market, the IMAPs carry a combined price tag of about $223,000—approximately $449,000 less than the channel banks. Credit for the rest of the decline goes to the elimination of DACSs, each costing $20,000 per location. The total savings of more than $469,000 in each of the company's 25 sites comes to roughly $11.7 million. Bear in mind that the windfall proved possible even though the provider never saw fit to concentrate its voice and data, often an IMAP's biggest source of financial benefit.

Summary

By installing IMAPs in their co-locations or local serving offices, the three subjects of the foregoing case histories benefited as in *Figure 7.*

FIGURE 7

Recap of How Three Carriers Gained Financially from Adopting an IMAP

Example	**Key Benefits**
Carrier #1	• Decreased capital expenses by more than $9.7 million over two years • Cut payments for leased backhaul lines by $45 million over two years
Carrier #2	• Reduced its operational costs by 61 percent • Trimmed its capital spending by 31 percent
Carrier #3	• Lopped 68 percent from its CAPEX (assumes channel bank @ $2,000/unit)

E-Business Management: A Practical Perspective

John McConnell
President
McConnell Associates

Introduction

E-business across the Internet continues to grow rapidly, and the challenge of managing these critical services grows even faster. The major challenges to effective management are as follows:

- Delivery chains are more complex, including multiple businesses, different types of service providers, back-end services, and content-delivery networks. It has been difficult to monitor, let alone manage, such a delivery chain.
- New opportunities and markets are being created by digital subscriber line (DSL), cable, wireless, and metropolitan Ethernet access. E-businesses must deliver total availability and superior performance as well as compelling content. Recent stock market activity indicates continued competitive pressures on e-businesses to make money as well as noise.
- The currently accepted tolerance point is eight seconds—the wait time for downloading a Web page or getting desired information—before a customer looks elsewhere. This tolerance will go lower as speed and performance become stronger differentiators. Slow e-business sites will be dropped for competitors that have compelling content with crisp performance.

The lack of management tools designed specifically to deal with e-business and its unique demands is one of the significant challenges. This report examines the e-business management challenge, defines a structure for evaluating various players, and compares some leading management products. An introductory overview sets the stage for a discussion of the segments of the e-business management marketplace.

An E-Business Overview

Segments

E-business has three primary segments: business-to-consumer (B2C), business-to-business (B2B), and Intranets—sometimes known as business-to-employee (B2E). B2C covers customer interactions with an e-business site; they may collect information, order products, and use support services through a Web site. B2C receives the lion's share of media attention, but B2B is where most expenditures have been targeted. B2B incorporates electronic processes, such as supply-chain management or customer-relationship management (CRM), as well as on-line services, such as trading or procurement. B2E supports an organization's internal processes; information is distributed through the Web and employees deal with more tasks on-line.

Many e-business transactions are actually hybrids. For example, a customer ordering a product at a B2C site may trigger B2B processes for verifying credit, preparing shipping notification, and initiating billing. B2E processes may also be activated to track inventory and distribute information about order volumes. An e-business manager needs to track this set of interrelated activities instead of strolling the aisles and noting activity. Traditional management tools are not able to cope with the complexity of most e-business environments.

Service-Level Agreements

E-businesses must have constant availability and superior performance to compete. Most are negotiating service-level agreements (SLAs) with their service providers, business partners, and suppliers to clearly define their expectations and to track delivery against objective service metrics. Service providers, partners, and suppliers also benefit by demonstrating they are delivering high service quality and identifying new opportunities for differentiated, high-margin services.

The Customer Experience

Another essential ingredient for e-business is tracking the customer experience. Compelling content, easy navigation, support, and performance affect the overall experience. A traditional manager had opportunities to interact with customers; e-business customers usually click to a competitor without bothering to complain. E-business management tools must provide insight into the customer's on-line experience.

E-business management solutions must deal not only with complexity, but also with the business that depends upon sophisticated technology.

Managing E-Business: It's the Transactions, Not the Technology

Any business is based on transactions: between a clerk and customer or between a customer and a Web site. For example, an on-line customer may query a Web site for product information, compare product features, select a product, provide billing and shipping information, and consummate the sale. The e-business usually responds with an e-mail confirmation and delivery details. Customers may return to the Web site to track their order, ask further questions, or change the selection.

Technology is the foundation, but it must be managed to optimize business transactions, rather than as an end in itself. For example, a network quality of service (QoS) solution will, at best, address only the network impacts on transactions. Network-management decisions must be guided by e-business goals.

Managing E-Business: It's the Business, Not the Technology

Managing an e-business means delivering meaningful information to the e-business managers, in addition to the technology teams. Business managers want information that allows them to understand what their on-line business is doing and where they need to focus their attention. For example, e-business managers are interested in business-related questions such as the following:

- Has the number of abandoned shopping carts changed after the Web page was improved?
- Are the new banner ads affecting our store traffic?
- What are my competitors doing? How do we compare?
- Are the inventory and shipping systems working effectively? Are we stocked properly? Are deliveries moving smoothly?

The e-business management system collects information from many sources and aggregates it into meaningful metrics for the business managers. It also provides detailed technical information to the technical teams.

An E-Business Management Architecture

Managing e-business transactions is an information and processing-intensive process. Information from many sources must be aggregated and correlated for a complete view of service flows. Large volumes of "raw" operational data must be converted to more meaningful information, such as baselines, revenues, and SLA compliance.

A simple architecture describes the e-business management system and its relationships to the technology infrastructures and their management tools (see *Figure 1*).

The dashed horizontal line is the boundary between e-business transactions and their supporting technology.

Above the Line: E-Business

The e-business management system deals with business-focused information such as the customer experience, transaction volumes, and supply-chain performance. It allows service providers and enterprises to offer differentiated services with the necessary guarantees. Customers are able to track their hosted e-business services and control their costs and service quality.

FIGURE 1

An E-Business Management Architectural Overview

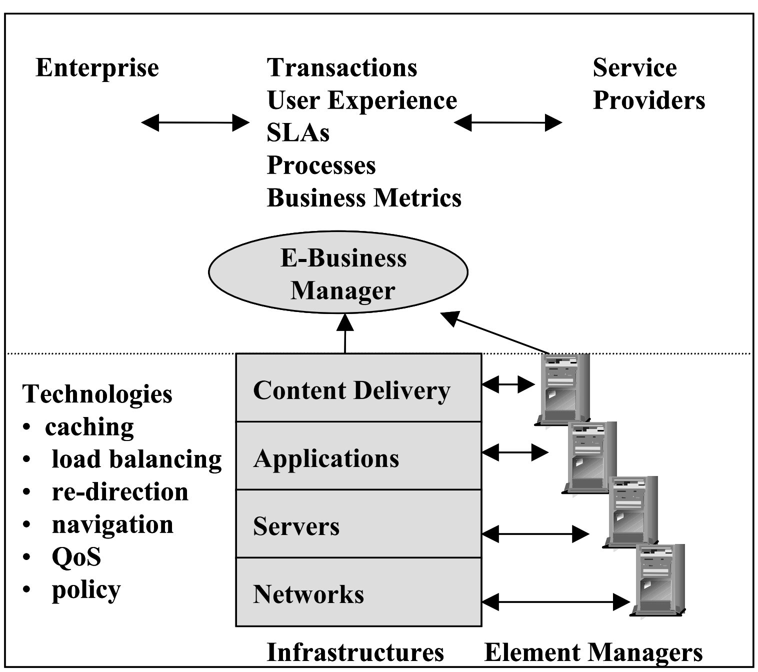

The e-business management system must also work with the technology-management systems to maintain and improve service delivery. Technology decisions must be informed and driven by their impact on customers, SLAs, and services.

Below the Line: Infrastructures (Segments)
Four technology infrastructures must be managed to provide high-quality transaction flows. Infrastructure elements are typically managed by vendor-specific element-management tools. Specific information from each infrastructure is used to assess its contribution to overall e-business performance.

Content Delivery
The content-delivery infrastructure manages the placement and optimal delivery of content. This model uses the content-delivery infrastructure to delineate services delivered by a service provider. Many of the same tools, such as re-direction and caching, are also used for the server infrastructure.

The complexities of managing content are continual. E-businesses constantly update their content to make it more attractive, to respond to new changes, or to offer special promotions, for instance. Creating and distributing information to the content-delivery infrastructure is a separate discipline from making certain that the content is distributed for fast access with load balancers, caches, or other devices.

Applications
The applications infrastructure is concerned with the functions needed for optimal application performance and service quality. The e-business management system must make certain that applications are operating correctly—starting them, halting them, or moving them to different servers.

The interactions between applications, content management, and monitoring grow more complex as dynamic (personalized) content becomes more common. Personalized, changing content strains the ability of the management system to generate the correct virtual transactions and to check the content when the transaction is completed.

Servers
The server infrastructure is focused on availability and high transaction volumes. Multiple sites and *n*-tier architectures are coupled with caching, re-direction, and load balancing for high-quality e-business services. The e-business management system must track transaction volumes and adjust the number of servers and task distribution accordingly.

Networks
The networks infrastructure spans enterprise and provider facilities and offers differentiated services for e-business transaction flows. E-business management is responsible for end-to-end performance even though external providers provide many of the networking services. Measuring performance and isolating problem areas is a major task.

Time Frames
Real-time monitoring, tracking, and adjustment are critical to maintaining the agreed-upon service quality. However, there are other management functions that occur over longer time intervals. Capacity planning, for example, is necessary for adequate and accurate provisioning. Resource shortages will cripple any site. Information collected by the e-business management system is used to forecast requirements. Other tools assess the impacts of new services or changes to avoid unexpected disruptions.

Defining a structure for discussing e-business management provides the foundation for comparing a range of available management products. The next section looks at specific needs and the capabilities of these products to meet them.

The Marketplace

E-business management is an emerging market, and many players want to position themselves in this new area. Some products are purpose-built for e-business management, while others are being adapted to address new market opportunities. Some target a portion of the problem, while others are growing a complete solution. The products compared in this report are shown in *Table 1*.

The products do not overlap completely, as can be seen. Firehunter, ProactiveInsight, and SiteScope were designed

TABLE 1

Vendor	*Product*	*Focus*	*Background*
Agilent	Firehunter	Full solution	Purpose-built
BMC	SiteAngel	Web application	Applications
Concord	Health	Application, server, network	Reporting
Freshwater	SiteScope	Web application, server	Purpose-built
Mercury Interactive	Topaz	Web application	Load testing
Micromuse	NetCool	Network	Event management
ProactiveNet	ProactiveInsight	Full solution	Purpose-built

for e-business management, while Concord, BMC, and Mercury Interactive entered with products originally designed for other management needs (reporting, application management, and load testing, respectively). SiteAngel is a subscription service, while the others are products.

Comparing products is always partially subjective. For example, user interfaces are evaluated on preferences, making them difficult to assess. A set of attributes was selected for this report to bring a perspective of current capabilities to the discussion.

Criteria

The criteria will be discussed briefly before they are applied to the products. They were selected to cover the important characteristics of a robust e-business management system. Each of the major categories may have several sub-categories. The selected criteria are as follows:

- ***Perspective and Breadth:*** The actual orientation of the product. Does it provide information on the business and customer experience?
- ***Monitoring and Instrumentation:*** Does the product collect information from a rich variety of sources with a choice of instrumentation techniques?
- ***Analysis:*** Does the product provide high-value analysis for tracking SLAs, baselines, and business metrics?
- ***Robustness and Scalability:*** Will the product have high availability and the ability to handle growth economically?

Perspective and Breadth

Perspective and breadth are important overall differentiators. An e-business management system must first of all have an e-business perspective: focusing on managing transactions, presenting business-oriented metrics, and tracking the customer experience.

Managers must be able to see transactions flowing across their e-business environments. They must also be able to correlate their services to the underlying infrastructures so they can understand the impacts of element problems on services, SLAs, and specific customers.

An e-business management system must have the breadth to cover all key infrastructures. End-to-end transaction-performance monitoring is essential. In addition, a solid e-business management solution also monitors the range of underlying services such as domain name servers (DNSs) or lightweight directory access protocol (LDAP).

Monitoring and Instrumentation

An e-business management system must track the behavior of elements, infrastructures, and e-business services. It must continually check for compliance to negotiated service-quality metrics and provide real-time alerts when situations require immediate attention from the management team.

Basic e-business monitoring must include transactions and sub-transactions (for higher granularity). Response time is an important metric, but content checking is becoming essential to insure the integrity of the transaction.

The e-business management system takes advantage of passive and active measurements. Passive measurements track actual activity, such as a desktop agent monitoring transactions with remote servers. Active monitoring generates virtual transactions that the management system uses for building consistent baselines.

Analysis

The collected data must be analyzed and presented to the management team. There are real-time and longer-term analysis tasks. For real-time analysis, data from many sources is organized to track the performance of various services and compare it to the SLAs. Both thresholds and baseline deviations are important indicators of potential problems.

Detecting a potential problem is only the initial step; further analysis is needed to identify the likely cause. This can extend from triage—identifying the infrastructure for deeper analysis—to identifying the element itself. Additional correlation is needed to associate elements with the customers, services, and SLAs that are impacted.

Longer-term analysis functions are needed for identifying usage trends and projecting future resource demands. Capacity planning prevents disruptions, because resources were not in place for growing e-business volumes.

Robustness and Scalability

E-business management systems must have high availability. Using multiple servers provides fault tolerance and failover when failures occur. Data management includes backing up key management information to protect it.

Scalability is also a concern because volumes can grow rapidly. Multiple servers also aid scaling. One of the most important areas is scaling the management staff. Highly automated management tools that discover the information that they need, adapt to changes, and relieve humans of mundane tasks are also needed for scalable solutions.

Results

Scoring

Simple scoring is used to evaluate relative capabilities. Satisfying a criterion earns a maximum score of 10 points; partial satisfaction scores lower. Weighting could be used to make specific criteria more important than others.

Perspective and Breadth

The specific features were as follows:

- Provisioning of e-business metrics
- Visibility of the entire service chain
- Visibility of the four infrastructures
- Measuring the user experience

Table 2 shows the scoring for these first features. It should be noted that differences are already apparent. As shown, several of these products do not address the basic requirements for e-business management. For example, BMC, Mercury Interactive, and Freshwater are not tracking all four of the e-business infrastructures. Customers do not have overall visibility and control of their transactions with a single product.

Some scores were also lowered due to incomplete coverage. ProactiveNet, for instance, only collects information with its own monitoring agents, reducing its visibility into the complete infrastructure.

There are many supporting services that must be monitored. For example, DNS or LDAP have substantial impacts if they are congested or unavailable. Automatically monitoring a wide array of services simplifies operations and provides valuable information (see *Table 3*).

SiteAngel and Topaz are currently focused on Web application and server performance; therefore, they do not monitor many of the underlying services supporting the Web transaction.

Monitoring
E-business management focuses on transactions, in contrast to the usual infrastructure monitoring tasks. Transaction monitoring must approximate the user experience by using virtual transactions—activity generated by the management system for management purposes (see *Table 4*).

Granularity is also a consideration. Complex transactions will activate sub-processes that are also important to measure. For example, if a credit authorization or a shipping/inventory search is slow, the customer satisfaction decreases, although the fault is externally based. E-business managers need the granularity to understand the behavior of their processes.

All products measure overall transaction performance, and most (except for Micromuse) also provide finer granularity. Thresholds can be set for sending alarms to signal the management system when performance degrades. E-business managers want alarms in close to real-time—some of this is dependent upon the agents. SiteAngel, in contrast, only provides periodic reports rather than alarms.

There is more separation when examining the four infrastructures supporting e-business. SiteAngel and SiteScope, in particular, only monitor the performance of the Web application and server, leaving the remaining infrastructures invisible.

SiteScope, Topaz, and SiteAngel are unable to provide business metrics because they measure the response times of only virtual transactions. They do not collect the needed

TABLE 2

Perspective and Breadth

Perspective	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury Interactive **Topaz**	Micro-muse **NetCool**	Proactive Networks **Insight**
	yes	no	yes	no	no	yes	yes
	yes	Web only	yes	Web & server	Web & server	partial	partial
content delivery	yes	no	yes	no	no	yes	yes
applications	yes	Web	yes	Web	Web	partial	partial
servers	yes	no	yes	yes	yes	yes	yes
networks	yes	no	yes	no	no	yes	yes
user experience	yes	yes	yes	yes	yes	yes	yes
Score	**70**	**20**	**70**	**30**	**30**	**60**	**60**

TABLE 3

Monitored Services

Services	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury Interactive **Topaz**	Micro-muse **NetCool**	Proactive Networks **Insight**
DNS	yes		yes	yes		yes	yes
DHCP	yes					yes	yes
HTTP	yes	yes	yes	yes	yes	yes	yes
HTTPS	yes	yes	yes	yes	yes	yes	yes
FTP	yes		yes	yes		yes	yes
LDAP	yes			yes		yes	yes
IMAP4	yes					yes	yes
POP3	yes		yes	yes		yes	yes
SMTP	yes		yes	yes		yes	yes
Radius	yes			yes		yes	yes
Score	100	20	60	70	20	100	100

information to track orders, revenues, or other critical business metrics.

Instrumentation

Active collection builds consistent baselines and provides continuous monitoring. This must be blended with passive instrumentation that collects information from real e-business flows—where the business-oriented metrics are derived.

Pervasive instrumentation is needed in a complex, distributed environment. Placing agents in different locations provides measurements across different parts of the underlying infrastructures. SiteAngel and SiteScope use only a single site to generate virtual transactions (see *Table 5)*. This is limiting because different parts of the networks infrastructure are not tested.

Openness is also a key; there are many information sources available. SiteAngel, SiteScope, and Topaz use only their own agents for information collection. While this may suffice for measuring overall performance, it leaves a substantial portion of the underlying infrastructures unmonitored, forcing customers to employ several monitoring tools and integrate the data themselves.

Intelligent monitoring is also essential: The management system must minimize network, server, and application loads while collecting quality information at the same time. SiteAngel, for example, adjusts the polling interval for greater granularity after a potential problem is detected. Firehunter also automatically baselines and allows for sustained thresholds—tracking an indicated problem for a time interval to determine if it is real or an artifact.

Analysis

The wealth of raw data must be analyzed and organized for effective service tracking, planning, and decision-making (see *Table 6*).

TABLE 4

Monitoring

Monitoring	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury Interactive **Topaz**	Micro-muse **NetCool**	Proactive Net **Insight**
Business Metrics	yes	no	yes	no	no	yes	yes
transactions	yes	yes	yes	yes	yes	??	yes
thresholds	yes	yes	yes	yes	yes	yes	yes
alarms	yes	no	yes	yes	yes	yes	yes
sub-transactions	yes	yes	yes	yes	yes	no	yes
thresholds	yes	yes	yes	yes	yes	no	yes
alarms	yes	no	yes	yes	yes	no	yes
Infrastructure							
servers	yes	no	yes	yes	yes	yes	yes
networks	yes	no	yes	no	no	yes	yes
firewalls	yes	no	yes	no	no	yes	yes
ext. services	yes	no	yes	no	no	yes	yes
Score	110	40	110	70	70	70	110

TABLE 5

Instrumentation

Instrument-ation	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury Interactive **Topaz**	Micro-muse **NetCool**	Proactive Net **Insight**
multi-site	yes	no	yes	no	yes	yes	yes
open	yes	no	yes	no	no	yes	yes
active	yes	yes	yes	yes	yes	yes	yes
passive	yes	no	yes	no	no	yes	yes
monitor end-end	yes	yes	yes	yes	yes	yes	yes
smart monitoring	baselines, sustained thresholds	change polling	change polling	no	no	thresholds	auto-baseline
Score	**80**	**40**	**60**	**20**	**30**	**55**	**60**

Tracking Compliance
One important feature is tracking compliance with SLAs. All products track simple SLA metrics, such as response time. However, e-business service quality is actually determined by a set of interrelated factors, such as network and content delays or load-balancing efficiency. Firehunter, NetCool, and ProactiveInsight separate from the rest with their compound SLA metrics.

Baselines
All products in this report generate service-performance baselines—indicating the envelope of normal transaction behavior. Content checking is another important distinction: The transaction must complete correctly, with the correct information in the pages. A transaction that completes quickly but does not deliver business value can be deceptive for the management team.

Root-Cause
Analyzing problems—and potential problems—is another critical area. The management team needs fast and accurate information that identifies the causes of degraded transaction performance. Triage is the first level of problem analysis—isolating the problem to one of the four infrastructures. Root-cause analysis identifies specific elements contributing to the problem and overlaps functions provided by element and infrastructure managers. Products such as SiteScope can identify problems but did not earn points because they do not cover all the infrastructures. Concord earns reduced points because they offer only some server analysis through its acquisition of Empire Technologies.

Scaling and Robustness
Management systems must scale easily, accommodating rapid growth and change, conserving scarce staff resources, and delivering bottom-line results (see *Table 7*). Factors considered for scaling include the following:

- ***Intelligent-Agent Management:*** treating agents as a group and modifying them all with a single operation
- ***Intelligent Agents:*** offloading servers and conserving bandwidth by local processing

TABLE 6

Analysis

Analysis	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury Interactive **Topaz**	Micro-muse **NetCool**	Proactive Net **Insight**
simple SLA	yes	yes	yes	yes	yes	yes	yes
compound SLA	yes	no	yes	no	no	yes	yes
service baselines	yes	yes	yes	yes	yes	yes	yes
content checking	yes	no	no	yes	yes	yes	yes
triage	yes	no	yes	no	no	yes	yes
root-cause analysis	yes	no	server	no	no	yes	yes
Infra-str correlation	yes	no	yes	no	no	yes	yes
Score	**70**	**20**	**50**	**30**	**30**	**70**	**70**

TABLE 7

Scaling and Robustness

Scaling	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury **Topaz**	Micro- **NetCool**	Proactive **Insight**
collector management	yes	yes	yes	n/a	yes	yes	yes
multiple servers	yes	yes	no	n/a	no		yes
intelligent agents	yes	no	n/a	no	no	no	yes
Robustness							
redundancy	yes	yes	no	no	yes	yes	yes
fail-over	no	yes	no	no	no	no	yes
Score	**40**	**40**	**10**	**0**	**20**	**20**	**50**

- ***Multiple Servers:*** scale even further and provide aggregation through hierarchy; fault tolerance is also enhanced with several servers

Fail-over mechanisms are also essential for high availability. ProactiveNet and BMC share an advantage in this area.

Summary

Agilent Firehunter has a narrow lead over ProactiveInsight, and the two enjoy a significant margin from the others considered in this paper. Both leaders showed consistent strength in each category, attesting to their balance and focus on the transactions of e-business (see *Table 8*).

The remainder of the field suffers from incomplete solutions or adapting other tools for e-business management. For example, Concord Health is a powerful reporting package that is adding transaction management through acquisitions such as FirstSense and Empire.

Many of the products covered in this paper are clearly not able to deal with complex e-business management demands. There is considerable activity centered on acquiring the missing parts through mergers and strategic partnerships. Those who are lagging will continue to evolve toward the model of the purpose-built solution. However, integrating a set of independent products is no easy task.

The highest scoring product, Firehunter, illustrates the value of a purpose-built product that addresses the e-business management challenge directly. Service providers, the enterprise, and their customers will all find higher value when they opt for an integrated solution.

TABLE 8

Summary

Summary	Agilent **Firehunter**	BMC **SiteAngel**	Concord **Health**	Freshwater **SiteScope**	Mercury **Topaz**	Micro- **NetCool**	Proactive **Insight**
Perspective	70	20	70	30	30	60	60
Monitoring	110	40	110	70	70	70	110
Instrument.	80	40	60	20	30	55	60
Services	100	20	60	70	20	100	100
Analysis	70	20	50	30	30	70	70
Scaling	40	40	10	0	20	20	50
Score	**470**	**180**	**360**	**220**	**200**	**375**	**450**

"Bandwidth Glut" or Strategies in Need of Repair?

Telecom Supply versus Demand in 2002 and Beyond

Russ McGuire
Chief Strategy Officer
TeleChoice

Mark Heaton
Optical Analyst
TeleChoice

Introduction

For the telecom industry, the "odyssey" of 2001 was nothing more than a relentless, stomach-wrenching drop to the very bottom of the proverbial barrel. In fact, *The Wall Street Journal* reported that the cumulative net earnings of NASDAQ companies between 1995 and the first half of 2000 were completely negated by the net earnings (i.e., losses) for these same companies between the second half of 2000 and the first seven months of 2001. As far as the financial markets are concerned, it is as if the technology boom of 1995 to 2000 never happened.

Many shell-shocked and wide-eyed industry veterans who managed to stumble out of 2001 are now facing the obvious question: What's in store for us in 2002 and beyond? For service providers, particularly on the long-haul front, this question is unfortunately framed by the market's troublesome and persistent theory of a supposed "bandwidth glut" throughout the U.S. market.

Clearly, if carriers hope to turn the tide of pessimism about their success in 2002, it is critical that they loudly, publicly, and without hesitation address the market's perceived "bandwidth glut" head-on. Is this current wave of pessimism warranted? Is the market really facing an oversupply of network capacity in 2002 and beyond? In a word, "no"—and for three pivotal reasons:

1. Installed dark fiber does not equal available capacity (*lit* fibers equal available capacity).
2. It costs far more to "light" fiber than it does to install it.
3. Carriers have no financial incentive to light new fibers without corresponding demand.

When it comes to analyzing supply and demand, the core question for many carriers to understand is precisely *where* in the market current demand levels outstrip supply. Likewise, it is equally critical for carriers to understand exactly which market segments have already been flooded with excess amounts of supply. Combined, these two extremes of the market help to guide informed carriers toward the market's most lucrative investment opportunities and away from situations with costly and difficult competitive dynamics.

TeleChoice has created a modeling tool for understanding the nature of telecom supply and demand and has performed exhaustive research and analysis to populate the model with information about the U.S. long-haul transport market. This tool, the Model for Advanced Capital Planning (MADCAP), has been created primarily to assist carriers in making strategic decisions about when, where, and how to invest capital in their network infrastructure.

TeleChoice has used the MADCAP tool to understand the current transport market and to examine the implications on the industry, using a number of different potential industry scenarios. This document discusses those findings, focusing on the current situation as well as three different potential industry scenarios. The report provides high-level guidance on how carriers and their technology vendors should approach the realities of today's market and how they should manage the opportunities and risks associated with potential industry outcomes.

What Bandwidth Glut?

As previously indicated, TeleChoice's recent analysis of the carrier industry in the U.S. market reveals that there simply is no widespread bandwidth glut. In fact, the analysis finds that more than 63 percent of the fiber routes between top U.S. cities are at or near full capacity. However, before diving too far into the conclusions, it is important to understand exactly what is included in this analysis.

In performing the analysis, TeleChoice developed a comprehensive network model that incorporated and analyzed data from all carriers that have lit capacity along key network routes. The model not only examines physical network assets (e.g., dark fiber), but also, more importantly, the amount of fiber that has been "lit" and is ready for service. Overlaid across this snapshot of route-by-route network capacity (i.e., supply) were the aggregated capacity requirements (i.e., demand) rolled up from individual network applications, again, on a route-by-route basis. The map in *Figure 1* illustrates the results from this overlay analysis of the entire industry's current supply and demand levels.

Once complete, the model pointed to many interesting conclusions, which at the highest level can best be summarized as follows:

- On less than 20 percent of routes analyzed, there is a current oversupply of lit capacity.
- Under a variety of industry growth scenarios, all analyzed routes will require significant capacity upgrades within the next several years.
- Several routes require significant upgrades immediately.
- Based on the volume of single-service traffic, several carriers could benefit immediately from higher density (e.g., 40 gigabits per second [Gbps]) transport systems.

In short, TeleChoice's route-by-route analysis of available capacity across the U.S. market reveals the bandwidth glut to be limited to a few routes but not widespread. There is a relatively healthy balance between current levels of supply and demand on most routes.

Growth Scenarios

Given the empirical evidence collected through the model, three representative growth scenarios have been developed to project supply and demand dynamics for the period 2001 to 2005:

"Doom & Gloom" Growth Scenario

This scenario (see *Figure 2*) assumes that capital markets remain closed to all telecom infrastructure activities, both long haul and local/metro. The impact of this scenario is that demand growth is constrained by the continued bottleneck in the metro and access portions of the network. Demand continues to grow at a healthy rate but not at the potentially explosive rates indicated in the other two scenarios. In this scenario, networks (long haul and local) are still upgraded to stay slightly ahead of demand.

"Free Money" Growth Scenario

This scenario (see *Figure* 3) assumes that capital markets quickly reopen for all telecom infrastructure activities. The impact of this scenario is that the local bottleneck is broken, resulting in tremendous growth in demand. However, long-haul carriers aggressively extend their networks by building new infrastructure into new routes and new territories. In this scenario, supply outpaces even the rapid demand growth in most segments of the network.

"Rational Expansion" Growth Scenario

This scenario (see *Figure 4*) assumes that capital markets are open for local telecom infrastructure activities but are relatively closed for long-haul infrastructure. Because of the investment in the last mile, this scenario results in tremendous demand growth. Long-haul networks must be aggressively upgraded to keep up with this demand growth. However, long-haul carriers stick to success-based growth on their existing routes rather than undertaking large new construction projects. In this scenario, long-haul networks are upgraded to stay slightly ahead of demand.

Regardless of how quickly the market grows, demand (particularly data) is expected to continue rising at double-digit and possibly triple-digit rates, thereby requiring carriers to continue "lighting" vast amounts of new capacity to meet this demand. We anticipate that the U.S. telecom industry will continue its historical practice of creating new supply as close in advance of increasing demand as possible without causing severe service backlogs.

Implications

Obviously, the telecom industry is currently under tremendous pressure. Stock prices for virtually every player in the industry are at record lows. Equipment vendors have seen

FIGURE 1

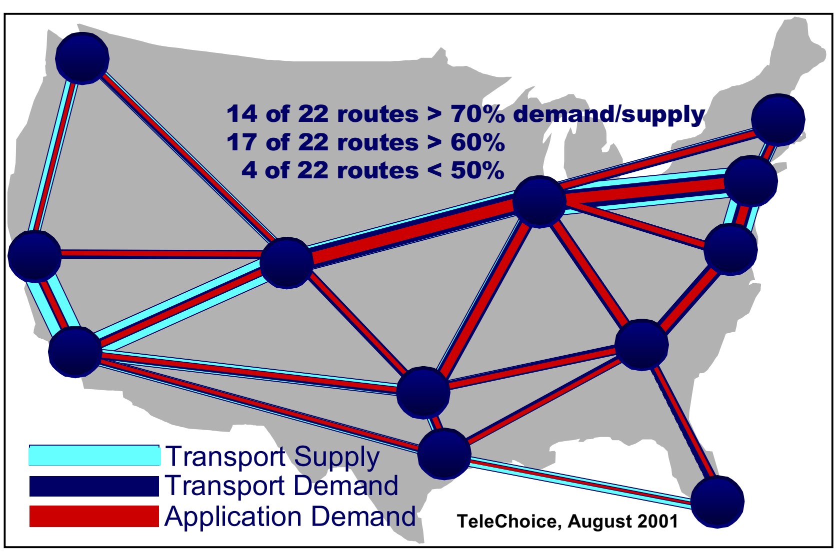

FIGURE 2

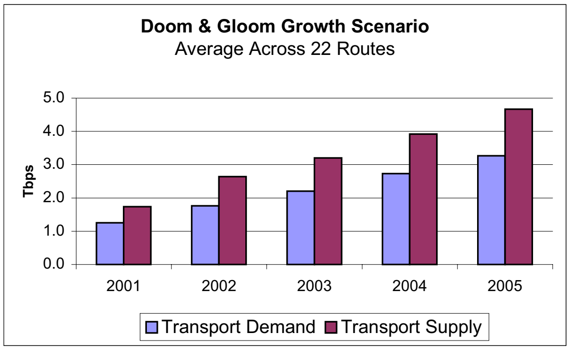

FIGURE 3

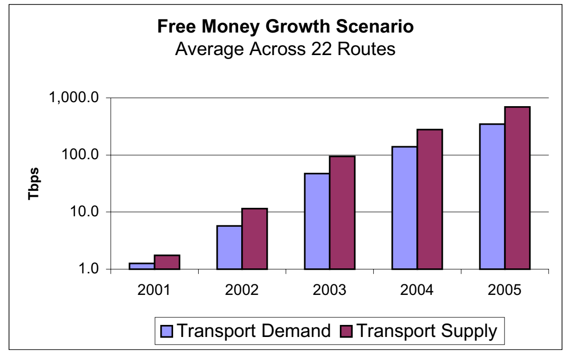

unprecedented revenue reductions. New service providers are in a race to achieve free cash-flow freedom before running out of capital while even many established service providers are struggling to work through the changing margin structures in the industry. Layoffs and bankruptcies have become a normal industry occurrence.

In the midst of this, how are we to interpret the analysis performed using the MADCAP model? The most compelling implications of this study can be summarized as follows:

- There is no widespread "bandwidth glut" in the U.S. telecom transport industry.
- The industry will continue to see healthy growth in bandwidth demand, even under the most conservative scenarios.
- Carriers must carefully consider when, where, and how to upgrade their networks to meet this demand. They must understand how their capital planning activities will play into the overall industry supply/demand picture, and they must carefully consider the implications of their technology choices on the scarce resources: primarily people and cash, but also floor space, power, and eventually dark fiber.
- Carriers must also realize that it is not an oversupply situation causing their pain. Rather, they must fix some fundamentally flawed aspects of their business model to build a business that can survive and thrive.
- Technology vendors must recognize that their falling revenues are not a result of a "bandwidth glut" and that other factors are impacting their business plans and sales forecasts.
- Technology vendors must carefully consider the technology capabilities that will fit the profile of the growing demand while addressing the pain being felt by their carrier customers. As they bring these capabilities to market, vendors must be prepared to encourage their customers to upgrade their networks while assuring them that the pain points (capital, people, space, power, fiber) will be well managed.

Industry-Wide Implications

Amidst this pain, the industry is struggling to figure out what has happened. A natural human part of that process, although not a healthy one, is to assign blame. Lately, the blame has fallen on the carriers that have spent billions of dollars to put fiber into the ground.

The implications if the carriers had failed to make this investment would be extremely painful for all of us. The network transport capacity that existed in 1994 would meet less than 2 percent of our current needs. The fiber construction projects and the technology innovations that have occurred since then have enabled us to enjoy a networking revolution that has greatly impacted our workplaces, our homes, and even the interaction between the two.

Furthermore, while pursuing those construction projects, the industry did not overbuild by putting too many fibers in the ground. The nature of network construction makes it very difficult and expensive to open the trench in the ground. Once the trench is open, carriers had better fill it in a way that they will not need to reopen it for at least the next 10 years, preferably 20. Approximately 22 percent of the fiber in the ground has been lit. Given that many of these projects have completed in the past year or two, this is a relatively high percent. Under the Doom & Gloom scenario, this will rise to 28 percent by 2005. Unless new technologies can more efficiently use the fiber, under the Rational Expansion and Free Money scenarios, all currently constructed fibers on the routes studied will have been consumed by 2004 to meet the projected growth in demand.

FIGURE 4

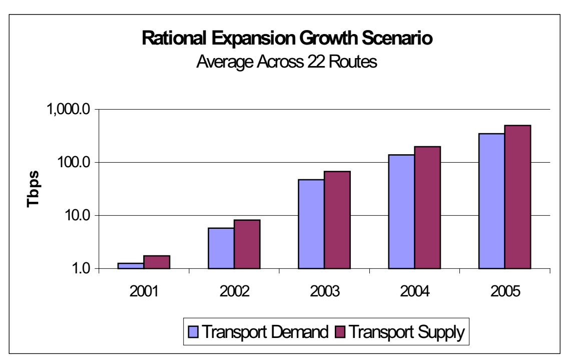

FIGURE 5

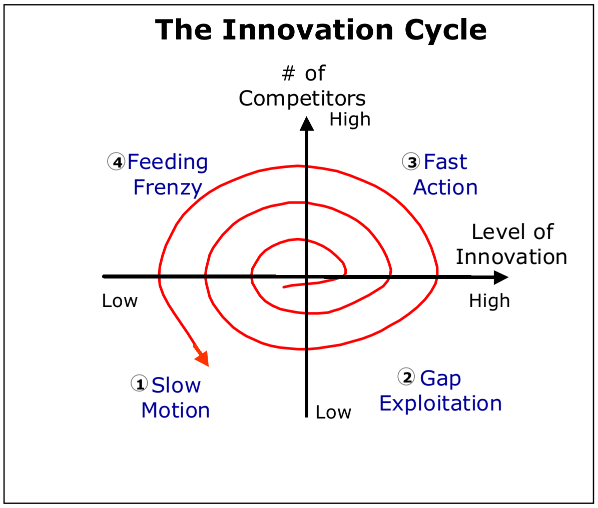

This study has also clearly shown that the industry has not lit that fiber in such a way as to create too much transport capacity, at least not on most of the major routes studied. Even on the four routes with too much lit capacity today, growing demand will eliminate that oversupply in one to four years.

However, what the industry did create were too many competitive carriers with too little innovation. Competition has become primarily price-focused, and as the industry works through the Feeding Frenzy phase of the Innovation Cycle (see *Figure 5*), the desperate players succumbing to the consolidation forces are leading the industry through irrational pricing and other irrational behaviors. TeleChoice believes that a number of clearly identifiable traits of companies are well positioned to survive and thrive in any phase of the innovation cycle. We have codified these as the "Seven Signs of Success," and we constantly work with industry players to develop these strategic traits within their businesses.

Meanwhile, on the equipment side of the story, as the industry worked through the Fast Action phase of the cycle, carriers were buying more equipment from more vendors than they knew what to do with, resulting in a tremendous inventory of undeployed equipment. At some carriers, the construction work in progress (CWIP) figure was as high as 40 percent of property plant and equipment (PP&E) even after their networks were largely completed. As companies of all sorts fail or become desperate for cash, additional equipment has become available at very low prices on the gray market. The net result is that equipment has been readily available for the needed network upgrades, without having to place a single order with an equipment vendor.

The good news for the industry is that demand growth of 20 percent to 722 percent will continue to create tremendous opportunities for all of the industry's survivors. Carriers need to work the irrationality out of their systems while they still have enough cash to carry them through. Equipment vendors need to wait as the oversupply bubble that currently exists in the market is worked out and, in the meantime, look for ways to drive sales of new technologies.

Implications for Carriers

The telecom carrier business is a scale business (see *Figure 6*). Transport services have become commodities with price as the primary basis for competition; new entrant carriers strive for the low-cost position. The technologies that they use to build their networks, combined with lean operations groups and new operations support systems (OSSs), have provided the basis for a lower-cost structure than the long-haul incumbents can offer. However, that low cost can only be achieved through high scale and high utilization. Only when the initial network cost and the corporate fixed costs can be spread across enough units will the cost per unit sold fall below that of competitors.

Most carriers have made fatal flaws in four critical strategic areas. While the industry and economy were booming, these flaws went unnoticed. Today, these strategy choices are causing incredible pain throughout the industry. The first area of challenge is in market strategy. Most carriers still model themselves after the historical model established by Ma Bell: To be "all things to all people." Most carriers lack any discipline for focusing on markets that they can uniquely serve. The result is that most carriers have an end goal of serving all telecom buyers (or at least all business telecom buyers) even though they may first focus on what is perceived to be low-hanging fruit. The second challenge is closely related: competitive strategy. Carriers generally fail to stop to identify their unique competencies upon which they can develop defensible differentiation. These two flawed strategies (or more accurately, missing strategies) have directly led to the price-focused market in which the industry operates today. The third challenge follows directly from the previous two: pricing strategy. The fourth critical strategy has been widely discussed throughout the industry by everyone with 20/20 hindsight: funding strategy.

The net result of these flawed strategies and the resulting economic realities has been a bitterly fought battle for every meaningful opportunity. Any customer representing a significant number of transport units represents a major opportunity to "push to the right" toward lower unit costs. Since pricing has become the primary competitive factor, prices have been pushed ever lower.

FIGURE 6

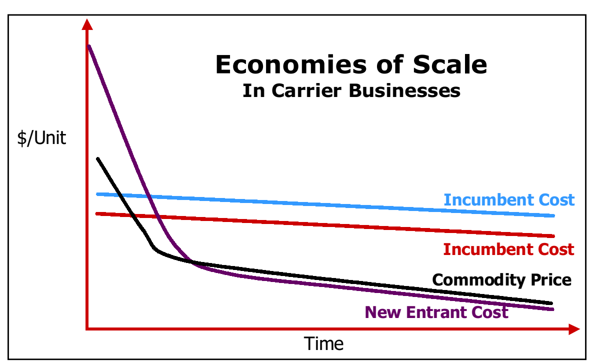

Compounding this problem, most of the new entrants have taken on significant debt in building their networks. Servicing that debt and meeting debt covenants has become an expensive challenge. Bringing cash in the door becomes ever more critical, even if just to push it back out to keep the creditors at bay.

So, in summary, the very real challenges facing carriers have not been caused by building too much network capacity, but rather by not building a well-tuned business:

- Focused on a winnable market opportunity defined by the company's unique strengths and assets
- Positioned and marketed to achieve compelling confidence from investors and customers
- Funded (and therefore structured/burdened) appropriately given the right business focus

Implications for Vendors

Technology vendors are similarly in a difficult position. For the past few years, vendors have been able to sell the industry more technology than the industry needed. Vendors increased their production capacities to meet this level of demand. Today, and for some period to come, sales volumes have not only fallen to a more reasonable level, but they have also shot way past "normal." Vendors' carrier customers are working through the oversupply inventory—both the additional equipment that they have purchased and the heavily discounted equipment available on the gray market.

Until this short-term oversupply gets worked out of the system, vendors will be sitting relatively idle. The greatest near-term opportunity for vendors will be to convince their customers that they need to purchase new technologies that are not in carrier inventories and cannot be found on the gray market. These new technologies must meet the carriers' current pains (capital efficiency, operational efficiency) while providing strong support for the coming growth in demand, both in terms of scale (based on raw volumes) and service attributes (based on the mix of applications and the unique demands of each service).

In short, vendors need to re-establish viable strategies around product innovation, operations scalability, and market creation. Vendors must re-establish discipline in these three key areas:

- Focus on the right product and market areas defined by the company's unique strengths and assets
- Collaboration with customers and other vendors to meet real needs with real solutions that leverage the capabilities and innovation that can be provided by the world leaders in multiple spaces
- Enabling service providers to offer profitable, non-commodity services that deliver true, measurable value to enterprise and consumer customers

Let's Improve the Market Performance of Our Joint Ventures and Alliances

Lana Phair-Sutherland
Principal
Phair-Sutherland Consulting

Internationalization, fast-developing technologies, and fast-changing markets are making alliances, joint ventures, and mergers more attractive and necessary for communications companies. There is statistical evidence that companies who do well at partnering have higher market values. At the same time, unsuccessful partnering can actually decrease market capitalization—and there is ample evidence that more than half of joint ventures and alliances fail to achieve their objectives.

This article focuses on improving the profitability of joint ventures and alliances through marketing and business strategies that are critical to success and are often underutilized. By pursuing strategies that are customer- and competitor-focused, these partnerships will enjoy positive revenue and stock performance.

Risks and Rewards for High-Growth Communications Companies

High-tech companies are at particular risk when entering into new relationships with high revenue goals, because, in the heat of pursuing a hot concept or a hot technology, the classic requirements of marketing warfare are often ignored. Yet marketing is critical to achieving those revenue goals.

A recent PriceWaterhouseCoopers "Trendsetter Barometer" finds that 30 percent of technology companies are using joint ventures to boost revenues and that more than 40 percent plan to use them in the coming year. Trendsetter chief executive officers (CEOs) achieve the following benefits:

- Acceleration of revenue growth
- Ability to increase profit margins
- Expansion to new domestic markets
- Sharing of scientists or professionals with unique skills
- Sharing of economic risk
- New product development
- Efficient commercialization of a technology or business concept
- Developing or acquiring marketing or distribution expertise
- Combining complementary R&D technologies

Five out of nine results sought are tied to effective marketing strategy, and the sixth, increasing profit margins, is closely aligned. Yet marketing issues invariably take a back seat in the formation and monitoring of joint ventures, alliances, and mergers. There are a host of synergies that may be gained by entering into new partnerships. However, neither cost savings nor rationalization will ever increase sales or profits.

The financial, governance, organization, and process issues that generally dominate the formation of various flavours of business partnerships are acknowledged as being extremely important and are very well dealt with in the literature and in practice. The focus here will be the market-related issues that drive profitability and performance.

Improving Market Performance in the Planning Stages

Post mortems on alliances and joint ventures that fail to perform show that partners and parents do not always take a proactive stance in considering areas where disagreements are bound to occur. Because marketing decisions are often more subjective than finance or technology decisions, disagreements in the go-to-market phase of the execution plan can cause frustration in personnel as well as ineffective market performance. *Figure 1* shows several areas that partners can address up front, in order to form guidelines for joint venture or alliance teams, once in the field.

Deciding on a Common Market Vision

Many joint ventures and alliances fail to meet the expectations of their parents or partners because they have been ill equipped to do so. It is critical, therefore, that the teams that do the up-front analysis focus on defining a defensible market vision and a practical approach to winning in that market. The assessments will include the following:

FIGURE 1

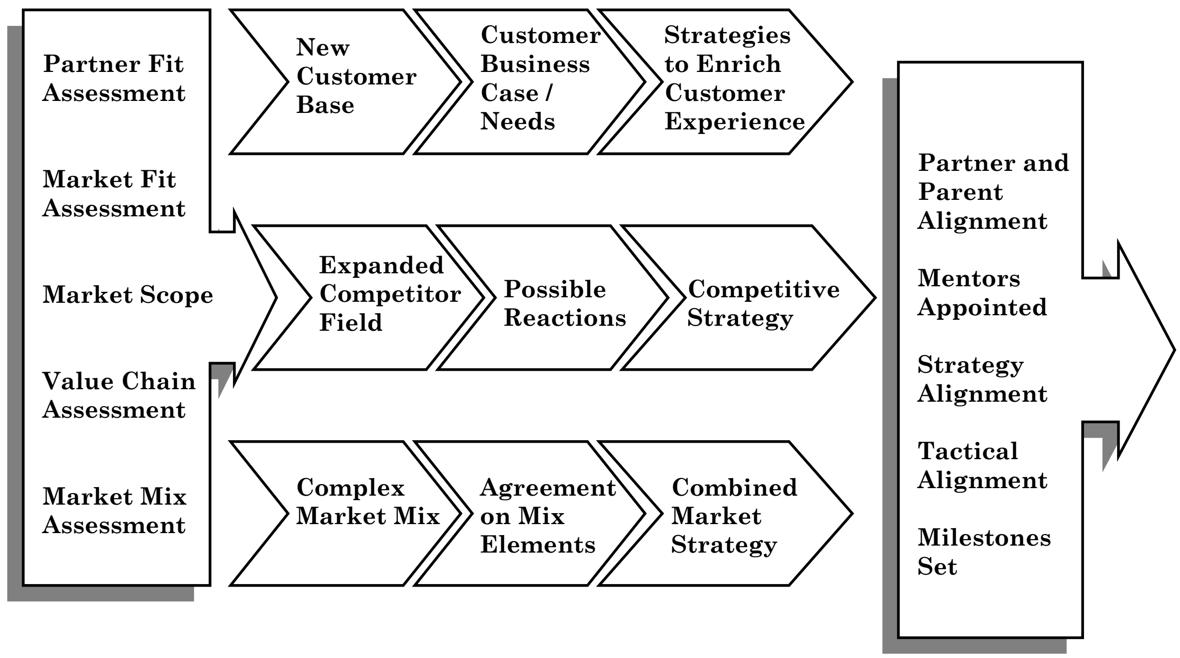

- The scope of the markets that the new entity can address: their size, growth, market windows, and barriers to entry.
- Customer profiles and business cases
- Competitive dynamics
- The nature of the value chain that the new entity will become part of and will use to satisfy its customers
- Product/market fit in light of the joint product portfolio

Customer Assessment
After the potential partners have assessed the fit of their operations and goals, their time horizons, and the financial returns that they hope to reap, it is time for them to turn their attention firmly toward the new base of customers that the new entity will enjoy. Beyond developing briefing formats so that each part of the new venture understands the business case or needs of the other's clients, the new partners need to do a full assessment of new customer segments and sectors that the partnership may open up.

- New customer groups should be defined and described.
- New customer groups need to be analysed carefully to ensure that their needs are fully understood and that there is a good fit with the benefits offered by the partnership.
- Traditional customer groups that will now be addressed by one or both partners as part of the new entity need to be revisited in terms of the best positioning for the new relationships that are addressing them, as well as for any new needs that can be fulfilled.
- Possible areas of conflict or current customer discomfort need to be envisioned and dealt with.

Once the market has been segmented and described anew, strategies can be developed to enhance the customer's perception of the joint venture or alliance and to enrich the customer's experience.

The Competition
Unfortunately, competitors rarely stand still in the face of market change. The first task of the assessment teams is to determine the nature of the direct and indirect competition that the new entity will face:

- Has the competitive field expanded?
- What are the major characteristics of the current competition? The new competitors? Major threats?
- Has the nature of competition now changed?
- What possible competitive reactions can the competition be expected to take?

Once the review of the competition has been conducted, in light of the product/market fit determined through customer analysis, a competitive strategy can be developed by the new team.

The Market Mix
With new customer groups and competition, and with a range of new internal and external players, the development of a marketing mix that will maximize revenue can be a complicated task.

Everyone knows about the 4 Ps of classical marketing, commonly referred to with a roll of the eyes as Marketing 101. However, it is surprising how many technology companies believe that Marketing = Communications. In these companies product decisions may be driven by technology and pricing by finance. Public relations and advertising may be outsourced. Sales and distribution, the customer facing groups, are often confused and ill equipped to provide a sustained product vision and service backup to the market.

In reality, Marketing = Optimizing Profit Performance. The importance of getting the 4 Ps of classical marketing right is not only in developing the optimal Product, Price, Place, or Promotion Strategy. The importance is in how the four are combined and executed.

Savvy partners strike teams to de-integrate the value chain and processes required to service each company's customers and re-integrate them in a competitive context in the new entity and the new marketplace. Pricing, Product, and Channel decisions in particular require careful attention up front. The following are examples of issues that might arise:

- Partners will often have historic pricing methodologies and pricing relationships with customers that may be challenged by the new relationship
- Companies may have cost structures that have historically dictated product price and volume decisions as well as sales force remuneration.
- Product lines that are targeted to current markets must be rigorously assessed to see if they should be offered in their entirety to new segments
- Can/should products be adapted? To what extent?
- Channels to market that have been in competition may now be common to the new entity. What channels have the most effective combination of cost and coverage to the new market segments? What product allocation and remuneration strategies must be developed to keep channels at peak performance?
- Poorly motivated sales and distribution teams are the undoing of many partnerships. Sales forces of the two entities may have been in competition in the past. They may still compete in other product-line areas. Remuneration structures are likely to be different. These issues must be solved and communicated well before any customer contact takes place.

The completeness and cohesiveness of the combined market strategy of the partnership will be one of the most important elements in its financial success. Strategy development requires some of the best minds from each organization, while implementation will require complete parent/partner alignment.

The Importance of Mentors

Many studies and the practical experience of joint-venture and alliance veterans indicate that success is often predicated on the ability of employees in the new entity to connect easily back to the partner or parent and to get things done. This need is important in all functional areas but is often critical in the marketing and sales department, with its need for targeted product development, timely shipments, and customer service.

It is highly recommended that a mentor in each partner company be appointed for each functional area manager. These mentors could come from the original assessment teams but should be in the same function as the new entity manager with whom they are paired, should be capable of "moving things along," and should be accountable for joint-venture and alliances milestones.

The triumvirate of mentors from each partner or parent and the functional area manager should set the milestones together and should also be prepared to monitor and react to changes that demand deviation from the original plan.

Improving Market Performance in Existing Joint Ventures and Alliances

Planning for optimal market performance will improve the revenue of new joint ventures and alliances. But how can existing partnerships that are under-performing benefit from these concepts? *Figure 2* shows a monitoring and feedback plan targeted to existing companies as well as to the ongoing success of new partnerships.

Counting on Change

One of the key attributes of the head of a successful joint venture or alliance is flexibility. Reporting to two companies while trying to drive a third company to profit targets is a demanding task, without the inevitable changes that occur in any market. One way of managing change is to plan for it.

Most companies perform to specific strategic milestones and tweak the direction of the ship as they go along. This works, as the only certainty about change is that it will occur. However, market and environmental change can be more detrimental to joint ventures and alliances because of their dual reporting structures and responsibilities. Parents or partners can have radically different responses to change and may themselves be undergoing changes that will impact the new entity. Finally, as we have seen, in technology companies marketing is not a science as it is in the consumer packaged goods industry, but rather a subjective and often piecemeal function.

For this reason, we advocate a formal change monitor function that is paired to quarterly reviews of milestones and that has a set of procedures attached to it.

Market Change

Markets are never static. Like all companies, joint ventures and alliances must face changes in industry structure, the regulatory climate, and economic conditions. Competitor's strategies are also likely to change, independently or in response to other change indicators.

The quarterly change monitor forces early recognition and assessment of the impact of change on the joint venture or alliance. In addition to parent or partner executives, mentors from each company are alerted to begin thinking about the role that each company will play to maintain a profitable course.

Customer Change

As we have seen recently, many telecommunications companies were unprepared for the swiftness and the level of change in orders from their customers and in the planning behaviour of their customers. While some of the suddenness can be laid at the door of just-in-time processes, it is also true that many vendors and operators did not look beyond their commercial relations to the problems that their customers were encountering. Some of the change indicators are as follows:

- Change in the customer's competitive playing field
- Change in the decision structure. Is it Finance,

FIGURE 2

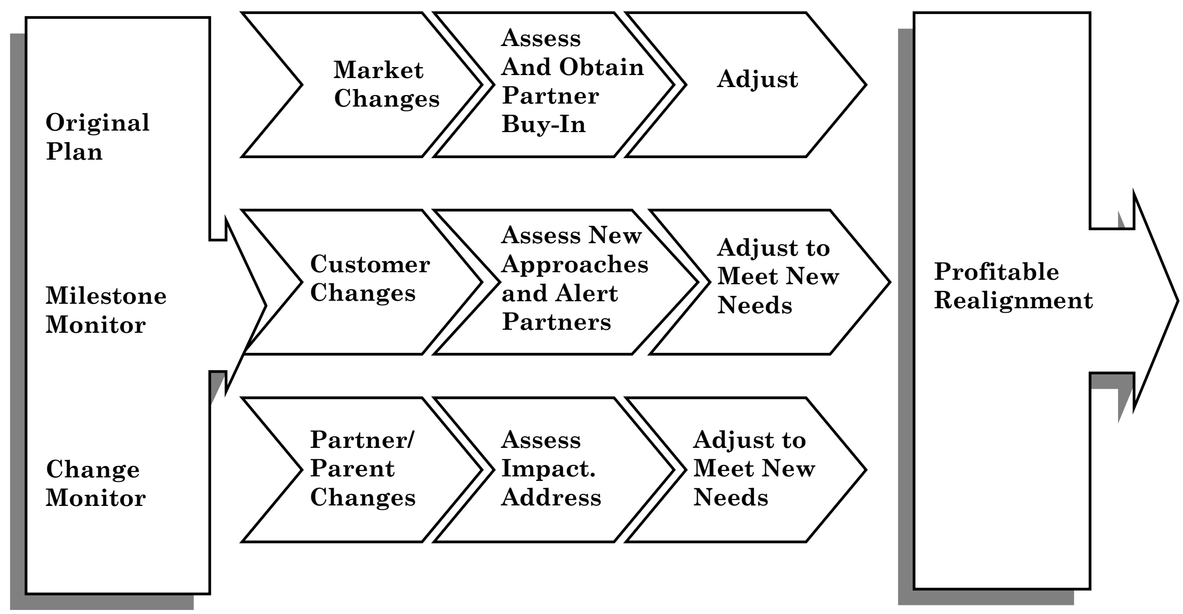

Technology, Marketing, or IT that is influencing buying decisions?
- Change in the economic viability of the customer's sector
- Change in market demand

When change monitoring is formally tied to quarterly milestones and is backed by the mentor structure, clarity is not obscured by fear of reporting potential problems, and adjustments can be made to manage costs and pipelines.

Partner/Parent Change
The objectives, resources, and organizational structures of partners and parents are prone to change, and these changes will impact the joint venture or alliance. Quarterly change monitor briefings and mentors for both companies should be used to assess and report on the impact that these changes will have on the ability of the joint venture/alliance to meet the needs of its customers and to operate effectively in its marketplace.

Summary

The structure and goals of joint ventures and alliances are diverse. Many are plagued with conflict potential from parents or partners that have been, or one day will be, competitors. Others are plagued with conflict potential from partners or parents that were caught up in the excitement of the initial deal but that did not put into place the processes and support mechanisms that would guarantee success. While joint ventures and alliances are becoming more numerous and more necessary in the fast-moving communications marketplace, and while more resources are allocated to them, a fairly low percentage of them meet their objectives. In addition, the employees who participate in them are often prone to extremely high levels of frustration.

Increased attention to marketing discipline in planning and operating alliances and joint ventures serves to focus participants objectively on customer needs and competitor realities. Establishing functional area mentors from each of the partner/parent firms, which advocate for the new entity and join in monitoring critical change areas as well as milestones, can improve both employee and company performance.

The Formulation and Implementation of Telecommunications Strategy

Steven R. Powell
Professor, CIS Department
California State Polytechnic University, Pomona

Abstract

Effective strategic planning requires a strategy that is both well formulated and well executed. Strategic telecommunications planning is no exception. This paper describes how the Balanced Scorecard approach to strategic management can make the formulation and implementation of telecommunications strategy more successful.

Introduction

Because of the turbulent and complex global environment in which companies nowadays must operate, current business practices emphasize speed, flexibility, adaptability, entrepreneurship, innovation, and employee empowerment (Muzyka 2000). Examples of such practices include rapidly deployable and sustainable competencies and skills, a phased experimental approach, dynamic and adaptive organization structures such as networks, and an environment where employees share in the value they create. Although the recent thinking in strategic management reflects these practices, the author has noted that relatively little of this thinking has been applied to telecommunications management (Powell 2001). In particular, issues that have gained significance in strategic management—such as vision, learning, and the role of competencies—have been neglected in telecommunications management.

Recently the author has become aware of the Balanced Scorecard approach to strategic management, which is now used by a variety of private sector, nonprofit, and government organizations, as well as the telecommunications departments within them (Kaplan and Norton 1996; Kaplan and Norton 2001). The Balanced Scorecard approach has enabled organizations that utilize it to more effectively implement their strategies, and when implemented in telecommunications departments, the approach addresses some of the strategic issues that traditionally have been overlooked. This paper examines how the Balanced Scorecard approach does this and how it can be used as part of a comprehensive methodology for formulating and implementing telecommunications strategy.

Strategic Management

Mintzberg, Ahlstrand, and Lampel (1998) have identified 10 schools of strategic management: the design school, the planning school, the positioning school, the entrepreneurial school, the cognitive school, the learning school, the power school, the cultural school, the environmental school, and the configuration school. The design, planning, and positioning schools are the older and more traditional schools of strategic-management thought. The seven younger schools better reflect current management practices. A brief description of each of the 10 schools follows:

1. The ***design school*** views strategy formulation as achieving a fit between internal strengths and weaknesses and external threats and opportunities. Direction comes from the top of the organization.
2. The ***planning school*** has the same view of strategy formulation as the design school, but formalizes the process, breaking it up into distinct steps. Strategic planners have a key role in the process.
3. In the ***positioning school***, generic strategies are followed in given industry situations. The planner becomes an analyst.
4. The ***entrepreneurial school*** views the formulation of strategy as a visionary process on the part of the chief executive, who becomes a creative leader.
5. The ***cognitive school*** looks on strategy formulation as a cognitive process, examining biases in the thought processes.
6. The ***learning school*** views strategies as being incrementally formulated and implemented by a process of experimentation. Strategists are found throughout the organization.
7. In the ***power school***, strategy making is rooted in power, either politically motivated within the organization or as power relationships among external organizations.
8. The ***cultural school*** views strategy formulation as a social process rooted in culture.
9. The ***environmental school*** looks at how organizations respond to particular environmental conditions.

10. The ***configuration school*** views organizations as clusters of characteristics and behaviors. Change comes from the transformation from one cluster or state to another.

Approaches taken in the formulation of strategy can cut across schools. Examples are dynamic capabilities (design/learning), resource-based theory (cultural/learning), chaos and evolutionary theory (learning/environment), and negotiated strategy (power/positioning.) All the schools agree on the following: Strategy concerns both the organization and its environment; strategy is complex; strategic decisions affect the overall welfare of the organization; and strategy involves issues of both content and process.

The Balanced Scorecard and the Strategy-Focused Organization

While formulating a good organizational strategy is certainly important, there is evidence that strategy implementation may be even more important, and to be implemented properly, the strategy must be measured and linked to the organization (Kaplan and Norton 2001). To accomplish this, Kaplan and Norton have developed a strategic-management framework based on the Balanced Scorecard approach, which translates an organization's mission and vision into a set of objectives and measures (Kaplan and Norton 1996). The result is a management continuum:

- ***Mission:*** Why we exist
- ***Core Values:*** What we believe in
- ***Vision:*** What we want to be
- ***Strategy:*** Our game plan
- ***Balanced Scorecard:*** Implementation and focus
- ***Strategic Initiatives:*** What we need to do
- ***Personal Objectives:*** What I need to do
- ***Strategic Outcomes:*** Satisfied shareholders, delighted customers, effective processes, motivated and prepared workforce

The Balanced Scorecard measures are balanced in several ways: external measures for shareholders and customers are balanced with internal organization measures. There is a balance between measures based on past results and measures related to future performance; and there is a balance between objective quantifiable measures and subjective judgmental measures. In particular, the Balanced Scorecard assumes that value creation comes from four different *perspectives*: (1) the financial perspective, (2) the customer perspective, (3) the internal-business-process perspective, and (4) the learning and growth perspective.

1. The ***financial perspective*** is the perspective seen by the shareholder, having as its objective increased revenues, improved cost and productivity, enhanced asset utilization, and reduced risk. The specific measures that are employed depend upon the organization, its industry, its competitive environment, and its stage in the life cycle. Many corporations use the same measures for their divisions and business units.

2. The objective of the ***customer perspective*** is to clearly identify the organization's target customers and business segments and to develop a core set of measures—such as market-share, retention, acquisition, satisfaction, and profitability—for the target customers and segments. Specific objectives and measures are drawn from three classes of attributes: product and service attributes (functionality, quality, and price), customer-relationship attributes (quality of purchasing experience and personal relationships), and image and reputation.

3. The objective of the ***internal-business-process perspective*** is the identification of the critical processes required to support the goals of the financial and customer perspectives. Three kinds of processes are incorporated in the internal-business-process perspective: innovation, operations, and post-sale service. In the innovation process, the characteristics of the market segments for future products and services are identified, and then the products and services offered to those identified segments are designed and developed. The operations process involves the identification of the cost, quality, time, and performance characteristics to deliver the required products and services. The last step in the internal-business-process perspective is the delivery of special services after the products and services are sold to the customers.

4. The ***learning and growth perspective*** recognizes that the organization must learn and grow in order to continue to fulfill the objectives of the financial, customer, and internal-business-process perspectives. There are three sources of learning and growth: employees, systems, and organizational alignment. The three core employee measurements are employee satisfaction, employee retention, and employee productivity. Because employees require accurate and timely information to be effective, there also is a need for excellent, capable information systems to support them. The last source of learning and growth, organizational alignment, deals with the organizational climate for motivation and initiative.

The Balanced Scorecard can be used as a tool for managing strategy in what Kaplan and Nolan refer to as *strategy-focused organizations*. Strategy-focused organizations are guided by five principles: (1) translate the strategy to operational terms, (2) align the organization to the strategy, (3) make strategy everyone's job, (4) make strategy a continual process, and (5) mobilize change through executive leadership.

This is accomplished in a four stage *strategic framework*: (1) clarifying and translating the vision and strategy, (2) communicating and linking, (3) planning and target setting, and (4) strategic feedback and learning.

1. In the ***clarifying and translating the vision and strategy*** stage, the senior executive team translates the organization's strategy into specific strategic objectives. Developing operational measures for the Balanced Scorecard enables vague notions of vision and strategy to be defined precisely so that consensus on the strategic objectives can be reached.

2. In the ***communicating and linking*** stage, higher-level strategic objectives and measures are disseminated throughout the organization so that employees can become aware of the factors most critical to the organization's success and can see how their role affects the entire organization. Higher-level strategic objectives

are translated into local measures at the lower levels. This process encourages a dialogue throughout the organization about long-term strategic, as well as short-term financial, objectives. Incentive compensation schemes based on the level of performance against the measures further ensure that organizational efforts are aligned with strategic objectives.

3. In the ***planning and target setting*** stage, goals for the Balanced Scorecard measures are set. Kaplan and Norton assert that the Balanced Scorecard approach is most effective when stretch targets are employed. Once the targets are set, resources are allocated by means of strategic initiatives aligned with the targets. This is accomplished in the Balanced Scorecard approach through a set of cause-and-effect relationships called *strategy maps*. In the planning and target setting stage, one-year milestones for each measure are forecast at the same time that the three-to-five year targets are set, thereby integrating the strategic and annual planning processes. In this way, the short-term targets serve as milestones along the long-term trajectory.

4. The ***strategic feedback and learning*** stage provides both a single-loop and double-loop learning process about the strategy. Single-loop learning results from feedback on how the organization is performing relative to its short-term targets along a long-term planned trajectory to fixed objectives. Double-loop learning recognizes that in a dynamic environment the objectives may change, in which case the trajectory must be modified. In the Balanced Scorecard approach, the assumptions and theory underlying the cause-and-effect relationships in the strategy maps are constantly checked with current observations and data. If after some pre-established time horizon satisfactory performance is not delivering the anticipated results, the quantitative relationship among the Balanced Scorecard measures may be changed or, if necessary, the strategy changed. Thus the strategic feedback and learning stage feeds back to the first stage, clarifying and translating the vision and strategy.

Traditional Strategic Telecommunications Management

As noted by the author previously, strategic telecommunications management has focused on the design and planning schools with very little attention paid to the other eight schools (Powell 2001). Strategic telecommunications planning supports the company's mission, is subsidiary to strategic business planning, and attempts to achieve a fit between (1) internal business requirements and organization and (2) external threats and opportunities in the technological, regulatory, vendor, and competitive environments. The process is formal, usually following the company's long-range and annual plan cycles. In addition to the required periodic reviews, ad hoc reviews may be conducted (Green 2001). The process consists of three distinct phases (Datapro 1986):

1. ***Review of Business and MIS Strategy and Their Effects on Telecommunications***
   The telecommunications strategy must not only support the company's business strategy, but also conform to its management information system (MIS) strategy, which also supports the business strategy. The key objective in this phase of the process is to analyze the company's business and organizational trends, and then to translate them into expected demand for services and products.

2. ***Projection of Environmental Trends and Effect on Service Levels and Costs***
   The objective of this phase is to analyze the key technology, regulatory, standards, and vendor trends and to forecast the impact these trends will have on the cost of telecommunications services and products and on the levels of service offered by the providers.

3. ***Evaluation of Alternative Methods of Satisfying the Communications Requirements in Terms of Cost and Service and Recommendation of a Set of Alternatives***
   The objective of the last phase is to evaluate alternative ways that the communications requirements can be satisfied and to recommend a particular set of alternatives. In addition to satisfying the given communications requirements, the chosen set of alternatives must meet established life-cycle cost, performance, and managerial (e.g., flexibility, service, support, and security) objectives (Ellis 1986).

Integration of the Balanced Scorecard Approach and the Traditional Strategic Telecommunications Planning Process

With its emphasis on balanced measures, strategic reviews, strategic initiatives, shared organizational vision, feedback and learning, and environmental analysis, the Balanced Scorecard approach incorporates a good deal of the recent strategic-management thinking described by Mintzberg et al (1998) and seems to address many of the deficiencies in the way strategic telecommunications planning has been performed traditionally. Although specific strategy maps and performance measures will vary from organization to organization, it is possible to draw some conclusions about how the Balanced Scorecard approach can be integrated with the traditional strategic telecommunications planning process based on the experiences of two current Balanced Scorecard users: the IT Department of Financial Service Company (FINCO) (Kaplan and Norton 2001) and the University of California at San Diego's Telecommunications Service unit of the Administrative Computing and Telecommunications (ACT) Department (UCSD Business Affairs 1998-2001; UCSD ACT 2000). When the Balanced Scorecard approach and the traditional strategic telecommunications planning process are consolidated, the management continuum becomes the following:

- ***Mission: Why we exist***
  The mission statement will emphasize the service nature of the organization. For example, UCSD's Telecommunications Services provides the campus with a range of telecommunications service, including telephone lines, cable installation and maintenance, short wave radios, pagers, cellular telephones, cable-television services, and data-network transmission and connection services.

- ***Core Values: What we believe in***
  The core values of most telecommunications organizations will be similar to UCSD's Telecommunications Services unit: the need for change, efficiency and cost-effectiveness, and strong customer relationships.

- ***Vision: What we want to be***
  The vision statement will express the telecommunications organization's desire for excellence so that the entire organization can be a leader in fulfilling its mission. Goals characteristic of this will be an effective and efficient use of resources, a high level of service quality, satisfied customers, reliable and innovative technology, and a high caliber of employees.

- ***Strategy: Our game plan***
  The steps in generating the strategy will be the same as in the past: (1) review the business and MIS strategy, determining their effects on telecommunications; (2) project the environmental trends and their effect on service levels and costs; and (3) evaluate alternative methods of satisfying the communications requirements in terms of cost and service and recommend a set of alternatives. For example, in their strategic plan, UCSD's Telecommunications Services unit considers infrastructure, data services, voice services, broadband services, wireless services, and emergency services separately, analyzing first the current campus situation, then campus and environmental trends, and finally future directions.

  Key to the mission, core values, vision, and strategy formulation phases in the planning process are the need to clarify the organization's overall vision and strategy, communicate and link the plan to it, and plan and set targets. This is accomplished using the Balanced Scorecard measures.

- ***Balanced Scorecard: Implementation and focus***
  The four Balanced Scorecard perspectives are the financial perspective, the customer perspective, the internal business process perspective, and the learning and growth perspective. As in the case of FINCO's information technology (IT) department, it is likely that the strategic objective from a financial perspective for the telecommunications organization will be a high return on spending, enabled by a wise amount of strategic spending and investment as well as effective cost-productivity measures. UCSD's Telecommunications Services unit measures profitability (percent excess revenue, accumulated earnings, and percent of net revenue), efficiency (total assets turnover), and leverage (total debt to total assets).

  The strategic objectives are not as clear-cut from a customer perspective as they are from a financial perspective. For example, the strategic objectives of FINCO''s IT department are to deliver (1) certain "basic," or expected, services reliably and at a reasonable cost and (2) certain "differentiated," or innovative, services that demonstrate intimacy with the customers' needs. UCSD's Telecommunications Services unit measures overall academic and administrative satisfaction in terms of its understanding of customer needs, accessibility via telephone, responsiveness, effectiveness of advice, resolution of problems, professionalism, courteousness, and movement in a positive direction.

  The strategic objectives from an internal business process perspective likely will be similar to those of FINCO's IT department: understand, anticipate, and prioritize customer needs; create and develop solutions; provide a flexible global infrastructure; manage technical and operating risk; and service the customer. UCSD's Telecommunications Services unit has used a variety of measures, including the following: minimum different log-ins and passwords needed by a staff member; monthly rate charged per user for one phone and one Internet protocol (IP) address; telecommunications cost per extension line; percentage of general assignment classes network connected as a fraction of total classrooms; and percentage of residence-hall connections as a fraction of total residence-hall beds.

  The strategic objectives from a learning and growth perspective should emphasize the development of capabilities. FINCO's IT department treats two kinds of objectives: employee and organizational alignment. The employee-related strategic objectives are to retain and train skilled people and to empower employees. Aligning and linking goals and rewards is the organizational alignment strategic objective. UCSD's Telecommunications Services unit uses a survey to assess organizational climate. Measures include the following: communication; compensation; customer service; decision-making; diversity; leadership; morale; performance management; teamwork; training and development; and vision, values, and mission. Interestingly, although both organizations deliver information, neither has a strategic objective dealing with its information system.

- ***Strategic Initiatives: What we need to do***
  The strategic initiatives are the specific programs needed to implement the strategy. For example, in its strategic plan, UCSD's Telecommunications Services unit lists a set of milestones and recent accomplishments, both for this year and for five years out, in the areas of infrastructure, data services, voice services, broadband services, wireless services, and emergency services.

- ***Personal Objectives: What I need to do***
  In the Balanced Scorecard approach, personal scorecards are based on a clear understanding of the strategic objectives for the entire organization, as well as those for the business unit, and are linked to the objectives. One way of accomplishing this is by means of variable incentive compensation arrangements tied to performance. The scorecard should be balanced, with objectives and measures distributed among all four perspectives.

- ***Strategic Outcomes: Satisfied shareholders, delighted customers, effective processes, motivated and prepared workforce***
  The strategic feedback and learning process in the Balanced Scorecard approach plays a key role in generating positive strategic outcomes. Without it, the strategic trajectory can miss its target (single-loop learning) or objectives may fail to be modified when

necessary (double-loop learning). Techniques such as correlation analysis, management gaming, and scenario analysis can be used to analyze the validity of the cause-and-effect relationships embodied in the strategy maps. A strategic-management review process, related to, but separate from, the operational-management review process, increases the effectiveness of the learning process. Strategic reviews generally are performed quarterly, while operational reviews are performed more frequently. At such strategic review meetings, strategic issues are reviewed, implications are discussed, and performance data is analyzed (Kaplan and Norton 2001).

Conclusions and Directions for Future Research

This paper has investigated the formulation and implementation of telecommunications strategy and has suggested how the Balanced Scorecard approach can complement the traditional way in which it is performed. The Balanced Scorecard approach emphasizes balanced measures, strategic reviews, strategic initiatives, shared organizational vision, and feedback and learning. These are important issues, which are addressed in strategic business management but are often ignored, or treated lightly, in strategic telecommunications management. Based on an analysis of how FINCO's IT department and UCSD's Telecommunications Services unit use the Balanced Scorecard, the paper presented a strategic telecommunications planning process synthesizing traditional methods and the Balanced Scorecard approach.

There are several directions for future research. As now constituted, the synthesized strategic telecommunications planning process is top-down in nature and quite formal. Recent strategic management thinking, as embodied in the learning school, views the strategy formulation and implementation process as an informal and experimental one. Reducing the formality of the process, and increasing the experimental nature of it, would be a welcome enhancement, especially for organizations with short time horizons. Another way that the process could be enhanced would be to increase the role that strategic capabilities, especially systems-related capabilities, play in it. Strategic capabilities were not discussed in great detail in the paper but are necessary for sustained growth. Finally, future research ought to be directed toward expanding the data base of telecommunications organizations to be studied. In addition to providing data to more fully address the roles of experimentation and strategic capabilities in the planning process, a richer data base will enable a wide variety of strategic telecommunications management issues to be investigated.

References

1. Datapro Research Corporation, "Strategic Planning in Telecommunication," *Datapro Management of Telecommuni-cations*, Delran, NJ, March 1986.
2. Ellis, Robert, *Designing Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1986.
3. Green, J. H., *The Irwin Handbook of Telecommunications Management*, 2nd ed., Irwin Professional Publishing, Chicago, IL, 1996.
4. Kaplan, Robert S., and Norton, David P., *The Balanced Scorecard: Translating Strategy Into Action*, HBS Press, Cambridge, MA, 1996.
5. Kaplan, Robert S., and Norton, David P., *The Strategy-Focused Organization: How Balanced Scorecard Companies Thrive in the New Business Environment*, HBS Press, Cambridge, 2001.
6. Mintzberg, H., Ahlstrand, B., and Lampel, J., *Strategy Safari: A Guided Tour Through The Wilds of Strategic Management*, The Free Press, New York, NY, 1998.
7. Musyka, Daniel, "Thriving on the Chaos of the Future," Mastering Management: Part One, *Financial Times*, October 2, 2000, pp. 2-4.
8. Powell, S. R., "The Strategic Management of Telecommunica-tions," *Annual Review of Communications*, vol. 54, International Engineering Consortium, Chicago, IL, 2001.
9. UCSD Business Affairs Performance Measurements, 1998-2001, http://www-vcba.ucsd.edu/performance/
10. UCSD ACT Telecommunications Strategic Plan, September 2000, http://www-act.ucsd.edu/telecom/planning/strategic/sp_view.html

Aspects to Consider When Creating a Fully Integrated Test Environment for a Telecommunications Carrier

Normand Primeau
Associate Director, Platforms and Architecture Lab,
Service Technology, Broadband Infrastructure
Bell Canada
Senior Member
Institute of Electrical and Electronics Engineers (IEEE)

Introduction

Major telecommunication companies, as other organizations, are facing day-to-day rapid technological changes. With the advent of fast technology improvement, telecommunications network elements (NEs) evolve and need to be verified and tested. These devices must be tested in captive environments to confirm their abilities to deliver the claimed vendor functions, to interoperate with existing platforms, and to verify their maintainability. To do so, test laboratories are implemented. In some cases, as need grows, the number of these testing facilities increases to a point where rationalization and efficiency must be envisaged.

Carriers' laboratories have materialized because there were drivers to implement them. Services based on new technologies and new concepts—such as asynchronous transfer mode (ATM), Internet protocol (IP), Gigabit Ethernet, etc.—are some of the recent triggers that could have required the building of a facility to perform interoperating testing. However, with time, some duplication of a different nature (equipment, facilities, mandates, etc.), representing misuse of efforts and funding, may have materialized. This is even more true in a large organization, because different groups, having very good intentions, would create these testing environments for excellent reasons, all aiming at serving legitimate causes, but independently.

Different departments—such as operations, technology development, service, support, etc.—may have established their own testing facilities to satisfy their needs without knowing if duplications existed. Obviously, any duplication must be eventually removed to gain efficiency. Gaining efficiency may mean a streamlining of capital requirements, sharing of resources, expenditures reduction, duplication of efforts and equipment reduction, leverage of subject-matter expert skills, etc.

The amalgamation of these testing facilities then becomes a challenge. Blending existing sites that sometimes are geographically disperse is not obvious. A plan and a strategy to do so must therefore be well thought out, and steps must be clearly established to ultimately achieve a world-class testing facility. Obviously, there must be more benefits to offset any inconvenient to do so.

Key aspects required to maintain efficiency in a laboratory/testing facility must be present. Process, test-results documentation and distribution, human resources, security, environment, etc. are important aspects that should not be neglected.

Operating a Laboratory Environment

Let us look at the process aspect. In fact, a testing facility must implement a process that will ultimately satisfy the organization's needs. The internal customers being served by these testing facilities must be aware of the process to fully benefit from it. Most of the time, the facility serves network architects who develop the technology to satisfy the service and product development needs (NTI/NPI new technology/product introduction).

As an example involving a technology-development testing facility, a process was established to filter testing demands. In fact, testing demands were originating from a multitude of different groups—e.g., service development, product development, operations, manufacturers, methods, customer systems engineers (CSE), sales, etc. It was a challenge to sort out which of the demands had the highest testing priority simply because each request was considered by the owner as the number one and had to have precedence over others. It became obvious that a filtering method had to be established, confirming that the testing requests were legitimate. In this case, "legitimate" means that a testing project being performed and which may last many weeks, if not months, should not be cancelled because the original trigger was not properly substantiated. This would therefore represent a loss of manpower and inefficient use of equipment. A way to avoid such scenarios was to use the network architects associated to each of the services to evaluate testing demands associated with those services. If actual technology or upcoming services would fulfill the needs, then the request would not reach the laboratory team. Also, if a similar request would reach the architect, it would therefore stimulate a standardization of a generic solution, and associated testing would be appropriately planned. Naturally, the network architects must be part of the service- and product-planning groups.

Other testing facilities may serve operational groups to satisfy various levels of support. In some instances, methods and procedures for the deployment of new technologies, or the upgrade of software bringing new features to existing platforms, must be verified to confirm efficiency and avoid a customer's loss of services.

Therefore, testing facilities can be used for all kinds of legitimate purposes, and the proliferation of such test environments is sometimes unavoidable. The question is how to apply such a principle to other testing facilities so that a unique approach is implemented. The answer will require careful attention. Implemented correctly, this aspect would serve to reduce non-valid requests and therefore mean more efficiency.

Testing Activity List

How does one track and control a laboratory output, knowing that unplanned activities regularly materialize? One would say, "Set up a list of projects and stick to it. Any new testing requirements would be added in a queue." This approach is uncertain. The purpose of such an activity list is to obtain an excellent view of what is the known workload to come, while maintaining some flexibility for the unknown. A list of activities must be considered as a tool to manage the manpower and the priorities and, at the end, satisfy all the laboratory customers.

Individuals who have been exposed to testing facilities know that religiously following a testing schedule is almost impossible. It must be very flexible and must be built to adapt quickly to important demands. Such a list can be called a "Living Activity List." It keeps changing as time goes. Moreover, if testing work is done meticulously, it is almost certain that issues and problems will be raised. Discovering problems and resolving them with manufacturers, or proposing alternate solutions, are all normal events that take time and argue against a strict and inflexible testing schedule. It is laboratory personnel's first role to unveil issues. This approach must be therefore maintained and encouraged.

Maintaining such a list is not an obvious task, as unplanned testing demands materialize without warning. As an example, a problem of generic nature occurring in the network would dictate its importance and become the top priority for testing, as it could affect customer's traffic and impact service. We are not talking here of day-to-day isolated operational problems, but rather of problems that can be associated with the deployment of new components or software that, in this case, impair an entire telecommunications platform.

Priority changes may also be triggered by special demands for end-to-end testing, a particular telecom solution for a large customer, which amalgamates a multitude of components never before associated together. They are sometimes called "individual case basis" (ICB).

The aforementioned are examples of testing demands that cannot be planned. Most of the time, these are important matters and therefore must be inserted into the work plan. A certain flexibility of human resources and equipment must therefore become an intrinsic part of the process. Managing a test laboratory involves constantly adapting to demands, reprioritizing projects, and re-assigning resources to satisfy needs.

From one quarter to another, the list is modified and activities are shuffled. However, one ongoing aspect is the recording of progress and of the completed testing activities. The completed activities are the mirror of the performed laboratory work. At the end of a calendar year, using such a list represents the total output of a testing facility.

Again, applying this principle to a unique virtual laboratory is not a given. Some type of similar tracking system must be envisaged. If a centralized group performs such a task, it would certainly permit each of the sub-functional testing environments to perform that at which they are good—that is, performing the associated testing work. Keeping some flexibility for unplanned requests, however, will become a challenge.

Test Plans

Various types of test plans exist. Test plans are used to identify what will be tested and how. These plans must be agreed upon with the originator of the demand to insure that the work to be done will satisfy the initial needs. Some are very detailed and can be large, reaching as many as a few hundred pages. Their nature or complexity depends on the maturity of the tested technology. The secret of using test plans is balance. In fact, some test plans detail each aspect to be tested, and the tester must identify whether the test works or not (pass or no pass), sometimes in a mechanical fashion. Whenever possible, when interoperability of equipment is being tested, a certain latitude must be left to the individual. It must cover the area that will be looked at. But at the same time, it must provide the liberty to investigate aspects that may not have been considered during test-plan write-up or must drill down unexpected behaviors discovered as testing unfolds. Most of issues and problems not found by manufacturers or discovered by telecommunications laboratories personnel originate from this type of testing flexibility.

Regrouping a multitude of sites into a unique virtual laboratory means that rules would have to be established and that potentially different type of test plans would have to be recommended.

Documenting the Results

Publicizing the test results, as well as their associated documentation and distribution, is also an important aspect in the existence of a testing environment. It represents the information on which the laboratory customers (service development, product development, methods of procedures) will base their decision-making. It can also be used by the operational departments to implement corrections and avoid future major service interruptions. It is also of importance for internal training purposes. Clear content, format, detailed data, summaries, etc. are also important. To satisfy different types of demands, a formal document format and a technical note format are required. When dealing with usual testing activities, which may last many weeks, the formal document template is used. However, when a punctual testing is urgently required, and results must be made available in a quick fashion, the technical note template is more appropriate.

Implementing the aforementioned aspect into a fully integrated virtual laboratory is one of the easiest aspects foreseen.

Information Storage

Storage of this information and its accessibility, while restricted to "need to know personnel" is also an important part of the success of a testing facility/laboratory. Any time spent on performing testing and discovering issues is worthless if not properly documented, distributed, and stored where the organization personnel can access the information. Having in mind an integrated test environment, a unique storage location must be implemented. With the intranet and Web-site event, access to such information is facilitated throughout an organization in a secure fashion when well implemented. This implies that all testing facilities follow the same rules and store the resulting testing information onto one easy but secure location. Again, integrating the storage of information for all associated sites of a virtual testing facility would be a straightforward activity.

Human Resources

Human-resources selection is the basis for the success of operating such testing facilities. Personnel with a variety of knowledge and experience are critical. When interoperability testing goes on, rarely one individual possesses all the knowledge to perform end-to-end testing. It therefore requires a number of individuals who together possess the necessary skills to do so. Laboratory personnel therefore require an acute ability to work as a team. It is critical to select types of individuals that will blend with the existing team. Laboratory work is very special and needs very special personnel. They need to be self-motivated, autonomous, curious, and eager to learn and, again, must work well as a team. Although a major aspect, technology expertise is not as important as it may appear at first. It can be easily developed through training and exposure to laboratory environment. One additional necessity is to be able to document the testing results. The entire laboratory process relies on the ability of the laboratory personnel to report the findings and clearly express the results. A balanced laboratory team must be formed from a blend of people with experience within the organization, in addition to new grads that are willing to learn.

When some individuals have acquired a large scope of testing abilities and knowledge, they usually reach an advisory role. One of the laboratory advisor roles is to ensure that the acquired expertise and methodology is passed on to new resources (coaching). The advisor also performs the revision of testing results documents. Among all kinds of other responsibilities, the advisor is also involved in performing in-depth testing of critical projects.

Challenges

The aforementioned are only some aspects that need to be considered when dealing with a testing environment/laboratory. Creating a fully integrated test environment for a carrier must include these aspects and more. It is even more challenging when a multitude of geographically disperse testing sites, originally having different functions or goals, must be virtually amalgamated. Amalgamating various testing sites and groups into one virtual laboratory would mean establishing a list of resources and their knowledge. This would permit a better use and leverage of subject-matter expert skills. Means of doing so include the implementation of a database or a way to share this type of information for better use of expertise.

Developing and implementing uniform methods of procedure throughout all laboratories has to be addressed. Most of the time, each testing facility performs activities that satisfy distinct groups, which have different mandates and deliverables. This aspect becomes critical as the control and maintenance of capital investment through a common inventory-management tool is established.

Linking different geographically disperse sites will help to achieve the optimization of equipment investment. A captive core network built using each existing location will then be possible. It provides an emulation of the production environment for all products and services. Ultimately, after physically linking these different locations, one should be able to perform testing from any location or part of this virtual laboratory environment and access this shared equipment. Remote testing of new network components located in one site will be possible to a certain extent. There will be cases where proximity to the equipment will be required. The question is up to what extent remote testing will be possible. Reality suggests that there will always be some required physical proximity to tested components for all kinds of reasons. Remote testing applies particularly well to software-feature and protocol testing involving small hardware changes. However, it hardly applies to physical-layer testing. Laboratory staff and engineers will have to adapt to such an environment. Reaching the highest level of remote testing will be a challenge that should be approached in phases. Process is the key.

Creating a fully integrated test environment will take time. The initial phases will be to establish equipment inventory and identify functions for each site. Subsequently, the integration of functions will have to be addressed. This will have to be performed keeping in mind the organization's common goals. A fully integrated test environment must also effectively serve the NTI, NPI, and operational verification office as a single entity. The management of such a virtual laboratory also involves an excellent understanding of budgetary needs (capital expenses) and, at the same time, the ensuring that the organization testing needs and priorities are satisfied.

This impacts the carrier's ability to deliver new services, new products, and, in other instances, the maintenance of their telecommunications network.

Conclusion

In conclusion, creating a single fully integrated test environment will be possible. Aspects such as budgeting, documentation rules, information storage, equipment inventory, the leveraging of subject-matter expert skills (lists), implementation of a captive core networks between sites, and some of the remote testing will be the easy parts.

The challenges will be to deal with the functions and the already existing responsibilities of each of the sites. Physical presence at the tested equipment will obviously remain for a good portion of the testing, at least initially. The presence of identical equipment at various sites will also remain to a certain extent, as it will be used simultaneously for different applications and purposes. This will become critical as the delivery of results to deliver new services and products will be required.

This will only be possible and will only bring benefits if it is done through well-organized steps.

Standards Progress and Industry Efforts Status

David Shrader
President and Master Talent
Master Consultant, Inc.

The year 2000 saw many battles in the intelligent network (IN) and Internet protocol (IP) industry in the standards and forum areas. Many new forums were created, many new standards were created, and some ended in unexpected ways. Some players were lost; some combatants decided to work together, which was a new development. Among others that normally work together, a major civil war arose. This paper reviews standards progress and industry efforts status.

TINA

The Telecommunications Information Networking Architecture Consortium (TINA-C) terminated its work, but its legacy lives on; all its specifications were officially completed. The Telecommunications Information Networking Architecture (TINA) remains very controversial. The TINA-C spent significant time working on business models and service-provider models. It provided specifications to write input into various different standards organizations in an attempt to influence the way those organizations worked. A number of these concepts still can be applied, but many believe the Internet won and TINA lost.

The Battle for Internet Telephony

A battle for Internet telephony (IT) took place recently. Session initiation protocol (SIP) and H.323 represent one of the most significant battles in this space. They continue to proceed independently, one in the International Telecommunications Union (ITU) and the other in the Internet Engineering Task Force (IETF), but they have begun to look more and more like each other.

As the ITU implemented H.323, it discovered that version 1 of H.323 required a long time to set up a call. When a person picked up the phone and said, "Hello," the other person did not necessarily hear it immediately and vice versa. That has changed. H.323 continues to take on more of the flavor of SIP and session description protocol (SDP). The fast start was introduced and version 4 of H.323 now exists, implementing more of the rapid creation of sessions that SIP provides.

On the SIP side in the Internet world, as service providers have become more involved in the deployment of SIP, enhancements necessary for a carrier network have been added, and it has become apparent that hooks are needed for quality of service (QoS) and services such as caller identification (CID), 911 service, and so on. As a result, many extensions have been added to SIP that make it resemble H.323. Although SIP and H.323 use different procedures and methods, they are evolving much closer to each other. The IETF now maintains a design team to address the interworking of these two protocols, developing documents and specifications on how that might operate.

The Battle for Media Gateway Control

Another interesting battle that has emerged deals with media gateway control (MEGACO). The IETF and the ITU worked together to define interfaces for MEGACO. In an unprecedented agreement, the ITU formed an alliance with IETF to develop the MEGACO protocol, which is jointly request for comment (RFC)–3015 and H.248, to define the interface between a gateway controller and the media gateways (MG)—e.g., the interface between the circuit and the packet networks. When the standard is considered, however, the final specification supports both the text version, which is what the IETF uses for all protocols, and the binary version. Thus, these two are not truly interoperable because one is in text and one is in binary. Other elements will be needed to convert from one to the other to be interoperable.

The Battle for Network Intelligence

In the area of network intelligence the ITU and the IETF continued to partition network intelligence protocols. The ITU study group 11 (SG11) extends the IN architecture to include access gateways from the IN to the Internet to interact with Internet servers. The IETF defines some of the actual protocols, however, that fit into that IN architecture. The IETF, for example, has defined the Public Switched Telephone Network/Internet Networking (PINT) protocol that allows the server on the Internet to send the query into the telecom network, which, for many users, is the IN, and allows it to invoke services on the IN side. This was completed in June 2000.

Interactions in the other direction are in the progress of definition in the Service in the PSTN/IN Requesting Internet Service (SPIRITS) working group. With SPIRITS, an IN device or telecom device will be able to access services and systems in the Internet world and, thereby, allow some sort of user customization of services by querying the user's computer directly or another Internet server. This work is still in progress. The PINT protocol is based entirely on SIP. The SPIRITS protocol is not sufficiently advanced to determine its base.

The Battle for Service Programmability

Once all the protocols and architectures are in place the issue becomes how to program services. The two major groups, despite the appearance of competition, actually have joined forces. The Parlay group was defining application programming interfaces (API) primarily for access from an untrusted space so that networks could be opened. The Java APIs for integrated networks (JAIN) Consortium had been working on Java-based APIs for service programmability, but needed the external, untrusted access. As a result, JAIN joined forces with the Parlay group, which benefits both organizations because the Parlay group now has a Java base in addition to its other, ongoing efforts. The Parlay group now has a buy-in in the Java-based community. The JAIN group now is able to address external service-provider access in a more uniform way. While Parlay API specification is introduced into the Java technology marketplace through JAIN, the JAIN initiative is increasing its reach by addressing the external service-provider access. The goal is to define a unified call model.

Protocols versus APIs

Another battle in the industry involves protocols versus APIs. The Internet model is built on protocols. In the standards group, everything is a protocol. Standardizing APIs is not encouraged within the IETF, although it has been done in the past. In programs, however, the thinking is in terms of APIs.

The TINA-C, which, as mentioned, has completed its work, advocated a fully distributed processing environment. Its entire structure was based on an API so that remote objects could be accessed using an API. Worldwide support of this, however, has not taken hold. While readily available distributed processing environments are seen as being possible within systems, within a particular vendor system, or within some closely coupled development environments, wide-open distributor processing environments are not yet common.

Parlay has taken drastic measures to promote acceptance by joining with JAIN and opening its membership. This, perhaps, will introduce the concept of APIs more and help build momentum. Many companies, on the other hand, are promoting the use of SIP to access applications without a distributed processing environment (DPE) or use of standardized APIs. Services can be bundled on a processor, perhaps using the extensible markup language (XML) method of coding those services. The thinking here is that routing the call using SIP and letting that processor handle it all on its own is inexpensive and easy; therefore, there is no need for letting distributed objects communicate with each other.

Instant Messaging and IN

Another new service that has exploded on the telephony scene and is trying to influence the IN community is instant messaging (IM). The issue is determining how the telco industry can take advantage of IM. A Presence and Availability Management (PAM) Forum has been created and focuses on this as a technology—how to provide it and how to make services that are based on it. As a result of this interest and effort, a civil war started in the IETF over how to provide IM and presence protocols. The IM and Presence working group in the IETF could not agree on one protocol. A set of requirements was agreed upon, but three protocols were proposed that satisfied the requirements for IM and presence. A stalemate resulted and no further progress could be made on one protocol. A group of experts was formed to identify the common elements and the differences in all the proposed protocols. This group defined a single model for transporting data. Then, because the differing elements were in the transport area, they defined three transport protocols, but all three transport protocols are for a common presence format designed to be interoperable via gateways.

ITU Telecommunication Standardization

The ITU started a new four-year study period. SG11, the lead group for IN standards, is developing standards for the SPIRITS capabilities and for how the IN invokes services on the Internet and takes advantage of what is available on the Internet to provide more advanced services or interaction with a customer. Study group 16 (SG16) in the ITU is developing multimedia communications standards. Formerly the home of H.323 and the other 320 family of standards, it is now the lead group for a new initiative called Mediacom 2004, which will hopefully define how multimedia services are done over the next few years.

IETF

In contrast, the IETF solves problems related to interoperability over the Internet. The one key group is the SIP working group, which currently is applying SIP to all aspects of interpersonal communications, from voice telecommunications, to voice calling to voice, to setup of multiplayer games, to videoconferences. It is also trying to do IM and presence transport for it. SIP provides a mechanism that will allow all that to occur over a transparent link. The goal is integrating all interpersonal communications.

The PINT working group has completed the definition of a protocol for Internet servers to access the IN. The management information base (MIB) is in final review. The SPIRITS working group is defining services for an IN element to query the Internet for information. The architecture is defined and protocol requirements are under development.

The IM and Presence working group of IETF has completed protocol requirements, and three subgroups are forming to define the IN transport protocols and the presence transport protocols. The MEGACO working group has completed H.248 base protocol. For every different type of device, however—the black phone, the computer, an integrated services digital network (ISDN) phone, and so on—a package is needed to support that device. The definition of supplemental packages is underway.

Mergers and Acquisitions: Making Them Work

Geoffrey Thomson
Business-Development Manager
Compaq Computer Corporation

Introduction

Mergers and acquisitions involve a host of complex interactions, and anyone involved with them will constantly face an ever-shifting set of challenges and questions. The more one is prepared for these challenges and armed with solutions to these questions, the smoother the transitions will go.

This paper discusses some of the challenges and issues involved with mergers and acquisitions, in an attempt to provide some possible solutions. Its purpose is not to define the one-and-only path to success, but rather to present some ideas, perspectives, and approaches that have worked in the past.

The Role of Change

Change is the one constant element in mergers and acquisitions. Change is inevitable in this arena, and although it can be stressful, it must be accepted. The many various groups who have an interest in any merger or acquisition are often referred to as “kingdom keepers,” because they tend to seek ways of locking their interests away in safety. This kind of resistance to change, however, can threaten the success of the endeavor.

Consolidation of the interests involved requires a rapid, concentrated effort to make sense of the elements in the merger or acquisition. A variety of company types, including vendors, carriers, and Internet service providers (ISP), will play roles in this kind of consolidation, and all these sectors must be addressed with tact and attention.

Unfortunately, the people responsible for making these transformations work are faced with limitations in both funding and time, and these limits can often cause them to make choices too quickly or without enough forethought. In the rush, seemingly obvious and important issues are often overlooked. And in an increasingly competitive environment, such oversight can lead to major problems. Companies who wish to survive in the current environment need to make sure they are examining their reasons for doing the things they do as closely and diligently as possible.

Convergence and Differentiation

It has been said recently that service providers today are all seemingly striving to become similar to one another. However, most service providers are, in fact, striving to differentiate their offerings in order to stand out from the competition. To do this, these companies must bring a number of different systems into alignment.

Such an alignment process can often be more complex than it first appears. Those companies that find themselves faced with only one or two vendors providing different applications and productsævoice mail systems, billing, and so onæare very fortunate; most companies have no such luck. And when multiple vendors come into play, the only way to sort through the maze of details and often-conflicting vendor requirements is to take the time to examine the situation in its entirety.

A converged solution may or may not be the best choice. In some circumstances, divergence presents a better option. The vendor community as a whole usually favors a converged solution, but it is not necessarily the right choice. No one can afford to throw away the parts of a system they are using on a daily basis, and those legacy items must be considered when selecting a solution. Choosing a single vendor in any area brings with it its own set of problems. In the rush to conclude a changeover, companies often forget to consider what is best for them in their particular operating environment.

A number of critical questions plague the evaluation process:

- When, if ever, does convergence present the most viable solution?
- How many of the individual solutions being considered are too expensive or too disruptive?
- Will these solutions, by their cost or their function, possibly injure the competitive posture of the company?

A Formula for Success

The first step toward success in the evaluation process is for the company to define itself and what it wants to do in both the short and long terms. Certain processes, procedures, and applications will show themselves to be critical to the company's business environment; and these cannot be eliminated. Once those critical elements are identified, the path forward becomes much clearer. Thus, this process of self-definition requires serious attention.

The second step, if the company plans to stay together as a unit, is to decide on one future application and use that as a guideline for choosing all the other solutions. Again, a clearly defined business plan is essential before taking any action.

Furthermore, it is highly advisable to bring in outside professional services to help plan, design, and manage this whole process. Too often, providers make the costly mistake of trying to do everything "in house." More often than not, though, the in-house approach fails miserably, as the aforementioned kingdom keepers bring all progress to a halt in their attempts to guard their territories. If they truly want to make the process work, companies need an independent, neutral party managing the change. The expense of bringing in these outside services will be easily forgotten when the process is successful, while any attempt to effect changes without such outside help is very likely to end in failure.

With regard to the solution itself, central to any viable option will be a sort of "traffic cop"æan element at the heart of the system that can pull everything together and manipulate and move the parts of the system around. The traffic-cop element enables a company to have a fairly complete view of the workings of the network and helps direct functionality and data movement. It provides the ability to move the critical data application solutions and to manage the network efficiently and effectively.

The Whole Picture

Sales, order entry, provisioning, and customer satisfaction are the critical components of a successful OSS solution. And if a company is looking after its network's health, it should be able to detect potential problem areas and see certain things almost before they happen. This kind of foresight and quick action make it much easier for the company to retain customers.

The key to success is to offer demonstrable, embedded value. If all the elements are managed correctly, the carrier or provider will survive. If not, the company will inevitably start losing customers and may find itself on the receiving end of a merger.

The choices made in mergers and acquisitions affect a company's agility, credibility, and profitability. Managed adeptly, these choices can move a company in the right direction and give it a strong competitive edge. A well thought-out and comprehensive OSS solution is usually the result of this kind of smart and savvy decision-making.

In going through these processes, the company should remember that asking for outside, professional help is not a sign of weakness, and that if they do not seek it out, their competitors probably already haveæand that puts the competitive edge on the other side.

MPLS and SONET Comparative Economics

Wael Yared
Principal
Cambridge Strategic Management Group

Synopsis

Multiprotocol label switching (MPLS) over a mesh network offers significant economic advantages, both in capital expenditure (CAPEX) and operating expenditure (OPEX), when compared with a traditional Internet protocol (IP)/asynchronous transfer mode (ATM)/synchronous optical network (SONET) network architecture. We model the comparative economics of these two approaches to estimate the potential cost advantage of MPLS, which derives both from less expensive equipment and from its more efficient use of network capacity. We also look at the basic economics of Gigabit Ethernet (GigE), which offers cost savings but generates lower revenues. Our analysis shows that a greenfield MPLS carrier leasing bandwidth and collocation space can anticipate investment payback in less than half the time required for a carrier using either an IP/ATM/SONET network or a Gigabit Ethernet network.

MPLS and SONET Economics

Cambridge Strategic Management Group (CSMG) is actively engaged in evaluating technology advances in communications, with a particular emphasis on the economic and financial implications of those advances. The following discussion focuses on MPLS and its ability to apply the economics of packet-switched networks to voice and other circuit-switched traffic.

Despite the recent downturn in many Internet-related businesses, Internet traffic growth continues to be extremely strong, running according to some estimates at up to 400% annual growth rate in the first quarter of 2000.[1] Data traffic in general is believed to have passed voice traffic in August 2000, yet data-generated revenue is estimated at only about 20% of total telecom revenue, with the rest derived from voice traffic. Cost-effective ways of carrying the rapidly growing share of data traffic are therefore of great importance to today's service providers, and they regularly assess new technology options that may yield lower costs. Conversely, on the revenue side, the ability to leverage the installed bandwidth base more effectively by carrying additional traffic with some flexibility is equally needed.

MPLS offers an attractive economic solution, while also providing critical technical advantages. It can work with any of the commonly used data-communications protocols (frame relay, ATM, IP) and provides a clear route to truly converged networking—networking that would require only one management level and one set of management tools. As a hybrid, sharing characteristics of both connectionless and connection-oriented networking, MPLS allows a traffic-engineering solution for IP, improving bandwidth utilization and delivering differentiated quality of service (QoS).

In the following analysis, we take the viewpoint of a new carrier, deploying its network without sunk costs. Our carrier will avoid having to acquire ducts and lay its own fiber—instead we assume that it will lease unprotected wavelengths as well as co-location space. Its only capital investments will thus relate to its MPLS network deployment. Our study examines MPLS's potential for applying the economics of packet data networks to the delivery of traffic, which today generates far higher margins—namely voice and video traffic. These applications require the traffic engineering currently lacking in IP alone. Using MPLS, however, a service provider can deliver the QoS assurances essential to competitive voice and video offerings, and can also offer virtual private network (VPN) and other managed data services.

For purposes of comparison, we also look at the economics of GigE, a simple protocol with very low costs, but without the traffic-engineering capabilities of MPLS.

We compare the current multilayered IP/ATM/SONET approach with that of MPLS and GigE, both from a limited technology perspective and from the perspective of revenues and costs. SONET offers a premium service with hard bandwidth guarantees, and typical prices for a digital signal (DS)–3 between major U.S. cities currently range from $500 to $750 per Mbps per month. For example, our research indicates that an interexchange carrier (IXC) can offer a DS–3 from New York to Los Angeles for about $650/Mbps/month. Such a private line can carry both voice and data traffic because of its dedicated bandwidth and built-in redundancy.

GigE, in contrast, offers service on a best-effort basis, with no guarantees. Today, GigE is confined to the metro area, and a very large metro connection (say 100 megabits per second [Mbps]) retails for about $300/Mbps/month. Because there are no quality guarantees, customers typically use such a connection for large database file transfers—a major pharmaceutical company, for example, might link its intra-city sites with a GigE connection.

MPLS, in contrast to both alternatives just described, can deliver a range of service levels:

- SONET–like guarantees at reasonably comparable (or slightly lower) prices
- An intermediate level of service aimed at less critical applications but still providing 100% packet delivery at significantly lower prices
- A best-effort service level like that of the open Internet, with no guarantees, at prices comparable to GigE. Traffic at this level is essentially pre-emptible, and no separate investment is allocated to it in our modeling

For example, a typical user of MPLS might deploy a number of 512 kilobits per second (Kbps) connections across different sites for videoconferencing at the highest service/price level, with e-mail and Internet traffic connections at the best-effort level.

MPLS thus offers considerable flexibility in service quality and a commensurate flexibility in pricing. The pricing flexibility derives from the lower capital and operating costs of basic MPLS equipment, and from the characteristics of the mesh network approach as compared with a typical SONET ring structure.

For illustrative purposes, CSMG has considered a carrier entering the inter-city market using MPLS and compares this carrier's situation with one using traditional SONET architecture. CSMG has built financial models of MPLS and SONET networks linking the top 15 U.S. markets. The SONET network is constructed of five interconnected rings, while the MPLS network is partially meshed so that each node is on at least three separate links (see *Figure 1*).

The MPLS carrier can enter the business with a very small upfront investment, because it can acquire unprotected wavelengths, co-location and last-mile connectivity using short-term contracts, and take advantage of quarterly to annual price declines in both connectivity and equipment. (It is essential, however, that the new carrier be able to lease at least three links out of each of the cities it serves.) Services of the three types listed can be offered on both the wholesale and retail levels. Wholesale customers include carriers such as competitive local-exchange carriers (CLECs), incumbent local-exchange carriers (ILECs), IXCs, and wireless providers, as well as Internet service providers (ISPs) and intra-city entrants such as metropolitan-are network (MAN) providers. Specialty applications providers are also a potential target: content delivery, storage services, and other application service providers (ASPs) can use the new carrier's network to connect to their customers as well as to connect their internal operations. Finally, large enterprises are an important target market for inter-city transport.

While many of these customer groups will focus exclusively on data, others (such as large enterprises) will also want voice capability. In nearly every case, whether or not voice or video is offered, service-level agreements (SLAs) will be essential to meet the demands of ultimate end users. MPLS can address this requirement while still offering price advantages and the option of best-effort delivery where appropriate.

Our financial model incorporates the costs for each of the required network components as well as the variation in revenue-generation between the networks, since SONET revenues will be higher due to its inherent QoS. We describe these elements in sequence below.

Capital Costs

Figure 2 compares the required equipment for the three implementations. The SONET network requires add/drop multiplexers (ADMs), digital access cross-connect systems (DACS), and ATM switches. The MPLS network requires transceivers, MPLS–enabled edge and core IP routers, and an optical cross-connect (OXC). For comparison, we also

FIGURE 1

SONET Ring and Partially Meshed MPLS Networks

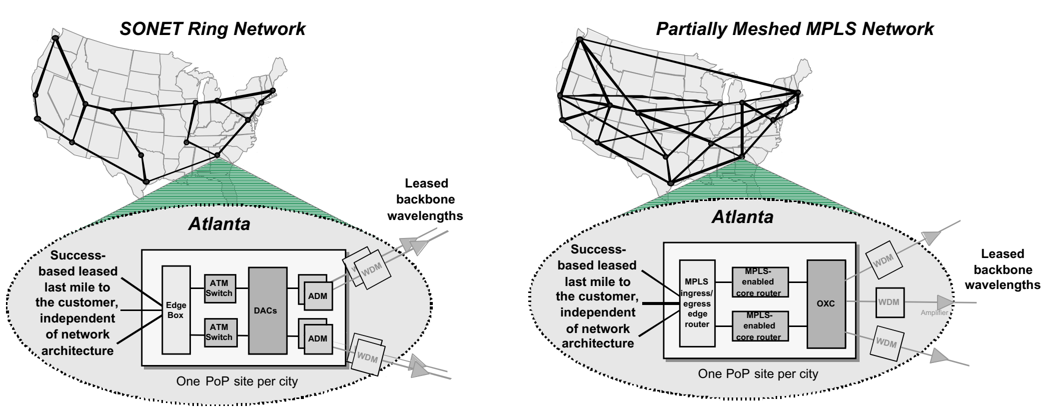

FIGURE 2

Network Components

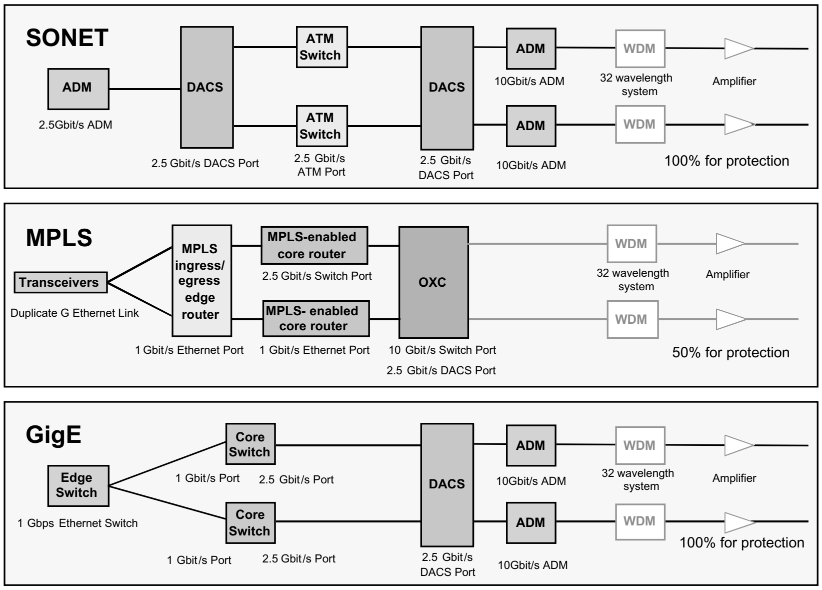

look at the GigE implementation. This requires an edge switch, core switches, DACS, and ADM (since our version is built on SONET).

We assume 50% protection for the MPLS mesh network, while SONET and GigE over SONET require 100% protection (1+1), as will be explained. The normalized effective cost per gigabit per second (Gbps), assuming 100% for the MPLS network, is 177% for SONET and 86% for GigE. Clearly MPLS offers a very attractive option based on investment costs alone.

We then look at the network-protection requirements for meshed versus ring topologies to quantify the bandwidth utilization rates in each deployment. Since ring networks have two links per node, and provide an extra full ring for protection, bandwidth utilization is static at 50%, regardless of network scale. In contrast, mesh networks have at least three links per node, yielding bandwidth utilization higher than 50%, and varying depending on the level of protection sought and the number of links at the node (more links means higher utilization levels for the same protection level). *Figure 3* makes the comparison, showing that the mesh network is always at least as efficient as SONET, and far more efficient under most topologies.

Another key factor in the cost of the network is how much transit traffic must be handled at each node. Traffic just passing through a node still requires bandwidth and equipment. In a ring, traffic typically passes through many nodes, since

FIGURE 3

Ring versus Mesh Bandwidth Utilization

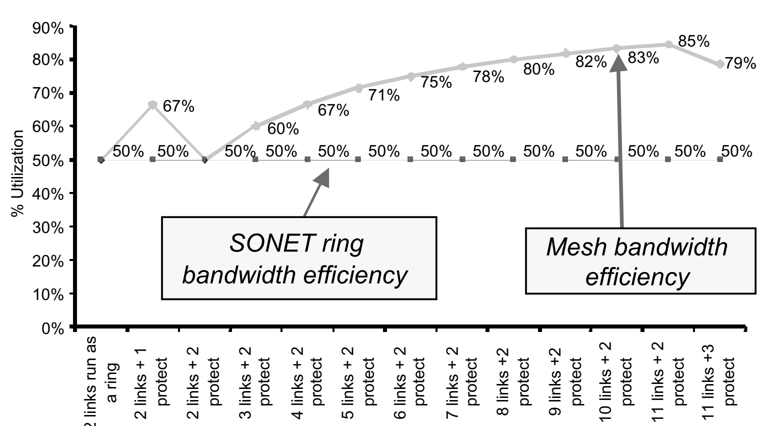

only adjacent nodes exchange traffic directly. In a mesh network, more cities are directly connected, and in a fully connected mesh, there is no overhead transit traffic. In a partially connected mesh, there is some transit overhead, but less than in a ring. For example, our sample analysis shows that our 15-city ring network would have to allocate an additional 279% capacity at each node for transit traffic. (We assume that all of the cities in our model generate similar levels of traffic.)

Combining these last two analytic approaches, we consider a representative connection that has 50 Gbps of destination traffic, using both SONET and an equivalent MPLS network. In addition to the destination traffic, the SONET network will have to handle a transit traffic overhead, which in our example we estimate at 75 Gbps. At similar levels of oversubscription for both networks in our comparison, this adds up to a total payload of 125 Gpbs and will require installed capacity of roughly 42 Gbps. In addition, however, a further 42 Gbps will be required for protection, equal to the capacity required for the actual payload. To carry the 50 Gbps of destination traffic with full protection, then, the SONET network needs capacity of about 84 Gbps or nearly 170% of destination traffic.

For a typical partially meshed network using MPLS, the same calculation looks quite different. Again, we start with the destination traffic of 50 Gbps, but the transit traffic overhead is far lower: in this case about 21 Gbps. To handle this total of 71 Gbps using similar oversubscription, nearly 24 Gbps of capacity is needed. To provide protection in a meshed network, only an additional 20% of capacity is required, bringing the required capacity to about 29 Gpbs or about 60% of destination traffic. (This analysis is summarized in *Figure 4*.)

FIGURE 4

SONET versus MPLS Comparison

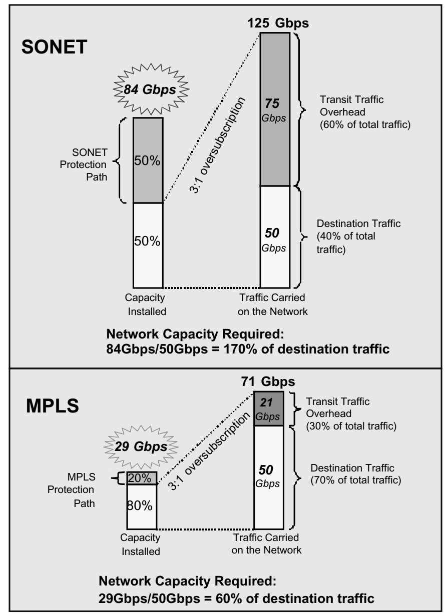

In summary, our analysis of capital costs shows three advantages for the MPLS mesh network when compared with a SONET ring, aggregating to a 50% savings in CAPEX when compared with SONET (see *Figure 5*):

FIGURE 5

Normalized CAPEX for Comparable Networks with the Same Traffic (Excluding Last Mile)

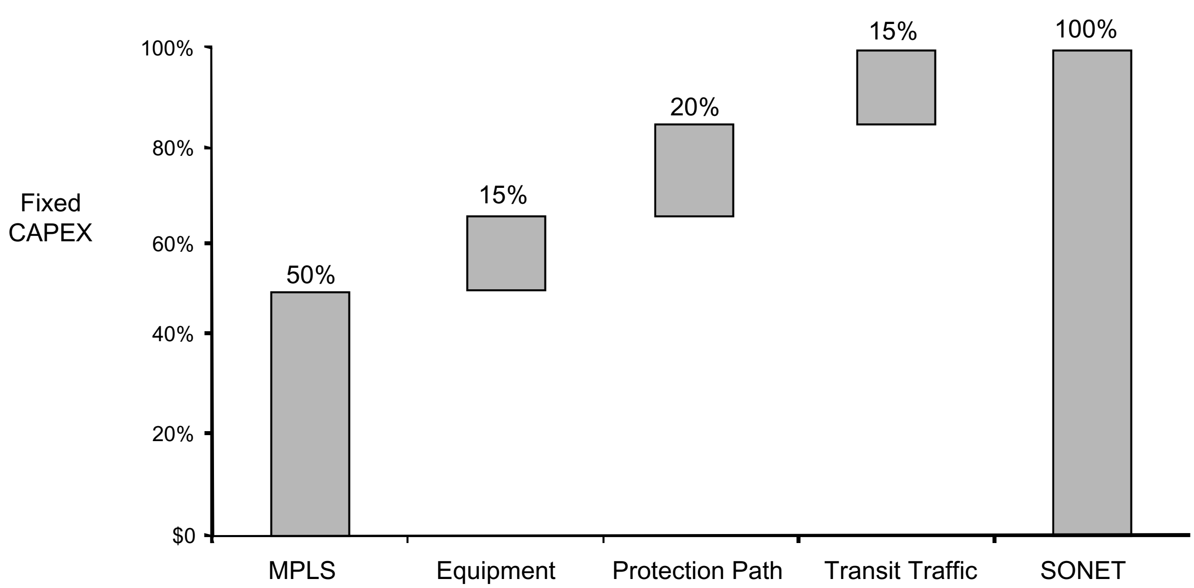

- Equipment costs are about 15% lower
  - A multilayer IP/ATM/SONET legacy network requires more expensive hardware
- Less protection bandwidth is necessary, saving another 20%
  - 50% of bandwidth is allocated for protection in a SONET ring, compared with a lower fraction for a mesh topology, resulting in the need for more inter-city bandwidth and higher equipment costs for the SONET implementation
- Transit traffic is handled more efficiently, yielding a further 15% savings
  - A SONET ring network requires increased capacity at each node to handle transit traffic

Operating Expense

Each of the three CAPEX advantages is mirrored in operating savings, and these account for a 43% advantage for MPLS. The savings, as indicated in *Figure 6*, come from three major sources:

- Network maintenance and upgrade costs are 18% of the savings
  - The greater complexity of the SONET leads to higher maintenance costs
  - Because the SONET is made up of rings, there is more inter-city capacity, resulting in more equipment that requires maintenance
- Equipment installation costs account for another 18% cost differential (We estimate that installation represents about 15% of the cost of equipment.)
  - MPLS requires less equipment and offers better network scalability
- Backbone leasing costs contribute 7% to the cost differential
  - MPLS uses backbone more efficiently, and therefore needs less backbone capacity

Financial Comparisons

If we consider the costs in conjunction with the revenue potential for the two networks, it is clear that MPLS will pay back its CAPEX far faster than SONET (assuming new builds in each case).

We assume revenue for MPLS of about 80% of SONET revenue, since the service does not offer the same level of quality. The lower cost of operating and upgrading a meshed MPLS network makes up for its lower revenue potential, so that MPLS and SONET have very comparable earning potential. The big difference is in the CAPEX, where MPLS has a significant advantage, leading to a more rapid payback.

Our financial model shows an EBITDA (earnings before interest, tax, depreciation, and amortization) margin of 42% for SONET and 52% for MPLS. MPLS's fixed CAPEX, as already noted, is about 50% of SONET's CAPEX, yielding a payback period for MPLS of just eight months, as compared with 20 months for SONET.

We can make a similar comparison with GigE. Revenues from a GigE deployment will be only about 38% of potential SONET revenues, due to its lack of service quality. Operating costs for GigE are similar to those of MPLS, meaning that GigE has only half the earning potential of MPLS. The capital savings for GigE over MPLS are modest: While MPLS

FIGURE 6

OPEX on a Fully Utilized Network

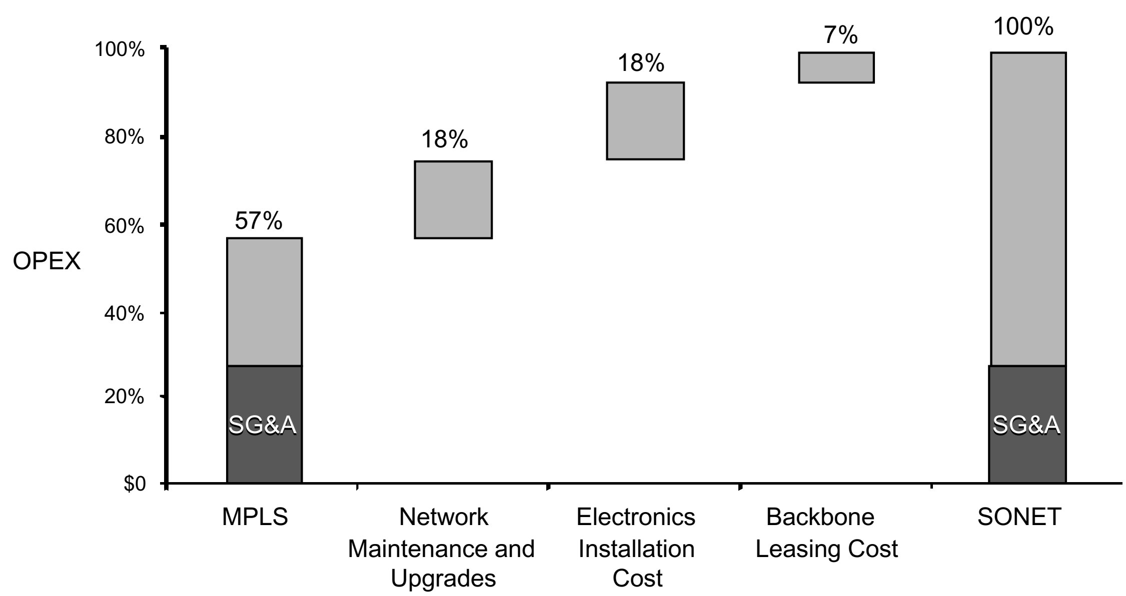

saves 50% over SONET, GigE saves 58%. As a result, an MPLS network pays for itself in less than half the time of a GigE network (18 months). *Figure 7* summarizes the payback periods.

FIGURE 7

Payback Period Comparison

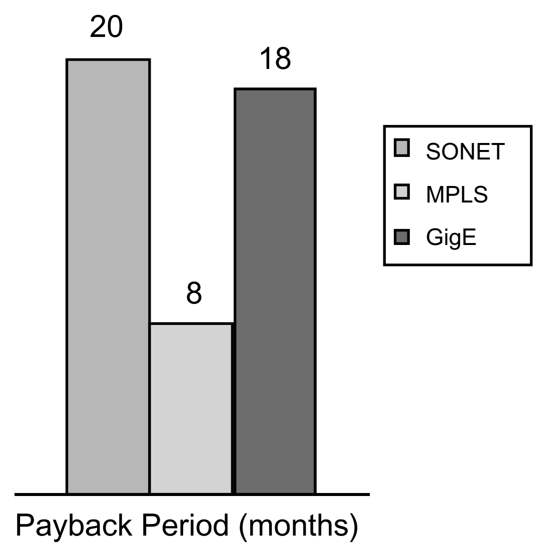

MPLS Trends

We believe MPLS will soon play a more important role in voice/data convergence, due to its theoretical capital and operating benefits. Like any new protocol, there continue to be technology and execution risks, which will slow initial acceptance. In particular, the QoS capabilities of MPLS have not yet been fully tested, and scale issues are not fully resolved, as there has been no full-scale deployment thus far. Interoperability, as with any protocol in an early stage, is also a concern. Also a risk is the tendency of proponents to oversell a new approach, leading to disappointment when it fails to solve every possible problem immediately and at low cost.

Nonetheless, some significant deployments occurred in 2001, including service offerings by Global One/Equant, Level 3, WorldCom/UUNet, Bell Canada/Teleglobe, Global Crossing, and AT&T. Trials were announced by Verizon, Qwest, Cable & Wireless, and BT, among others.

Looking ahead, we see three broad approaches to traffic engineering among today's major service providers. The first, which we term the "Brute Force Method," solves redundancy and quality issues by over-provisioning the network. If bandwidth is cheap enough, traffic engineering becomes an unnecessary expense. Carriers with this option, typically GigE MANs and ISPs, are unlikely to implement MPLS in the near future. The second approach treats different traffic flows separately, depending on the time requirements of specific types of traffic. Carriers using this approach—such as ILECs and public telecommunications operators (PTOs)—are usually technically conservative and will not migrate to a new network technology without well-documented performance data. The third approach is a hybrid one in which service providers experiment with MPLS to assess its operational benefits. They will allocate network capacity within a single network to different traffic classes and will use traffic engineering to reduce costs. IXCs are likely to be in this group, as are very large enterprises that manage their own communications networks and that must compete for internal resources (both dollars and technical personnel). It is from the organizations using a hybrid approach that most significant near-term MPLS deployment can be expected to come.

Notes

1. Caspian Networks, *US Internet IP Traffic Growth*, August 2001.

Section III:

Network Technology and Applications

Solving the Remote–DSL Challenge

Mark Abrams

Director, Product-Line Management

Advanced Fibre Communications

This paper considers solving the remote digital subscriber line (DSL) challenge in three different ways: converging edge applications driving the new network, providing universal broadband in the context of actual demographics and geographical issues, and integrated multiservice access platform (IMAP) solutions to meet the full sweep of demand in a profitable manner.

Converging Edge Applications

In the beginning there were voice services at the edge: 64 kilobits (k) plain old telephone service (POTS) and data services up to T1 speeds. Eventually an optical network was put in the core over a circuit-based network. In reality, quality of service (QoS) was the element providing revenue. The phone always worked, and the T1 had variable data hit rates. Customers could count on reliable services they purchased as consumers, and the service providers focused on building networks that accommodated these kinds of requirements.

Over the past 20 years there has been an explosion of data services. The network has accelerated all kinds of services being offered at the edge. This is being done primarily in packets. It is fundamentally a data service that is being offered, with new services added through the delivery of data alone. To accommodate this growth, the optical core continues to grow with optical carrier (OC)–12 types of services. The key is now revenue through bandwidth. With DSL, consumers will pay more for higher speeds of rates—T1s, T3s, and so on.

Further progress, however, will require the integration of these services. In the long run, the integration of voice and data services at the edge of the network is what the industry is striving to accomplish. At the same time, the optical core is growing. OC–48s and OC–192s are beginning to be deployed, and the issue is revenue from integrating variable bandwidth with QoS. What is needed is the ability to put both voice and data services on the same network and to use quality service mechanisms to generate revenue appropriate to the services being offered in the most cost-effective way. IMAP products will be a way for these solutions to be found throughout the network. The end result of this convergence is illustrated in *Figure 1*.

There are more and more digital carriers in the network. In 1999 11.9 million lines were added, and digital loop carriers (DLC) accounted for 88 percent of these lines (see *Figure 2*). An increasing number of services are being deployed over DLC. As a result, the challenge is to find ways to meet the needs for deploying these advanced services over digital carriers. The industry must solve this DLC challenge in the most effective way to maximize revenues at the edge of the network.

Today's Urban Broadband Technology

Digital subscriber line access multiplexers (DSLAM) have made an impressive start. They can serve 12,000 to 15,000 feet, and an extensive telecommunications infrastructure already exists to enable this service to be deployed quickly and easily. In addition, the demographics are highly predictable. For a town of a given size, it is relatively easy to determine who is there, what the chances are to sell specific services, and how to market to that specific group.

Furthermore, high-density populations mean cost-effectiveness for service delivery. A large building, or central office (CO), usually exists with all the necessary space, power, and environmental control to drive virtually any type of telecom or datacom equipment. This is a very low-risk proposition for carriers thanks to the law of economies of scale. In putting a 500-line DSLAM in a CO in a town that serves 100,000 people, the demographics and the population to a large extent ensure that the DSLAM can be filled up and, as a result, the optimal price point can be achieved. Thus, DSLAMs have worked well as DSL has been rolled out throughout the network COs.

The Remote Challenge

However, CO DSLAM solutions address only about 10 percent of the total geographical area that they actually cover. The area beyond 12,000 feet represents an area that can be 800 percent to 1,500 percent greater than the CO radius. A great deal of territory has lower population densities, and new construction and new customers are difficult to predict. The end result is that network planning and service delivery are significantly more difficult. The outcome, then, is what might be called broadband-deprived areas. When customers call to order DSL service, they find that they are served by a DLC and therefore unable to receive these kinds of services. Solving this challenge will be critical for continued growth in the industry.

There are three alternatives for deploying broadband using telephony techniques, though CO–based technology cannot

FIGURE 1

Converging Edge Applications

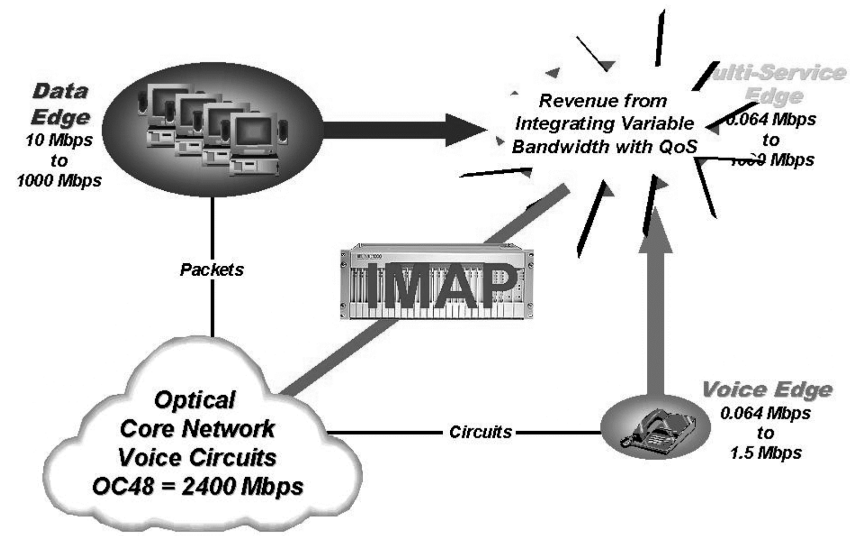

deliver beyond the CO. This leaves only two choices stand out in terms of extending the broadband reach beyond the CO. One, discrete box-on-box systems, puts a remote DSLAM next to the remote terminal. The other, broadband multiservice access platforms, integrates DLC functions along with the DSL types of solutions.

Box-on-box solutions are one way for the networks to build out in many places. Products can be mixed and matched to meet the service needs as they arise in a given area. This works well in COs where a multiplexer might carry the services out and low-speed shelves could add POTS, specials, and integrated services digital networks (ISDN). A digital system cross-connect (DSX) with some T1s in the area, or a T1 repeater bay, might be added to provide high-capacity T1 services at that site. Or these services might be provided over high-bit rate digital subscriber line (HDSL), which can be added to leverage the capabilities of that service. The newer synchronous HDSL could perhaps be used to take advantage of symmetrical DSL and asymmetric digital subscriber line (ADSL) for business and residential services. The result is a series of platforms that may or may not be well integrated, but that can meet all service needs as growth occurs (see *Figure 3*).

FIGURE 2

Overarching Deployment Trends

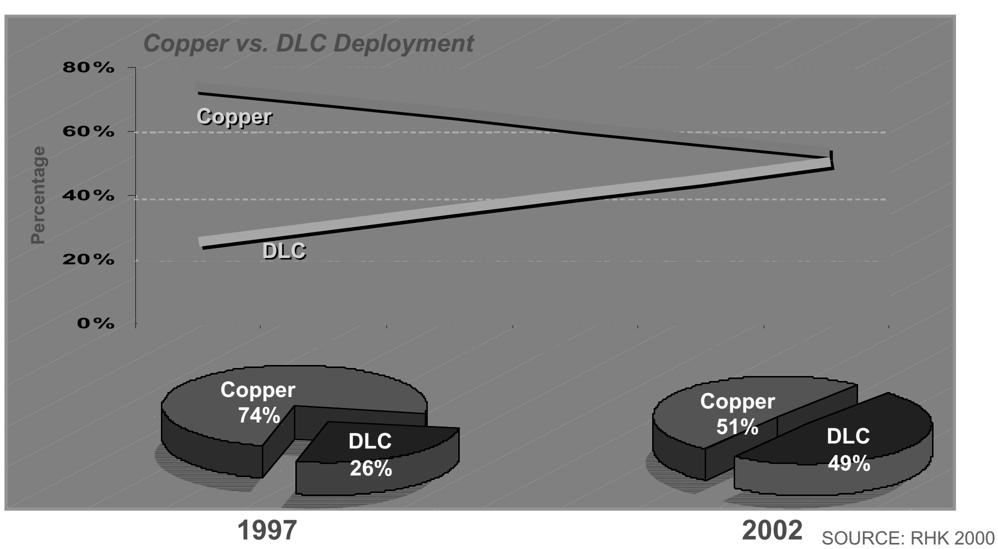

FIGURE 3

Box-on-Box Solutions

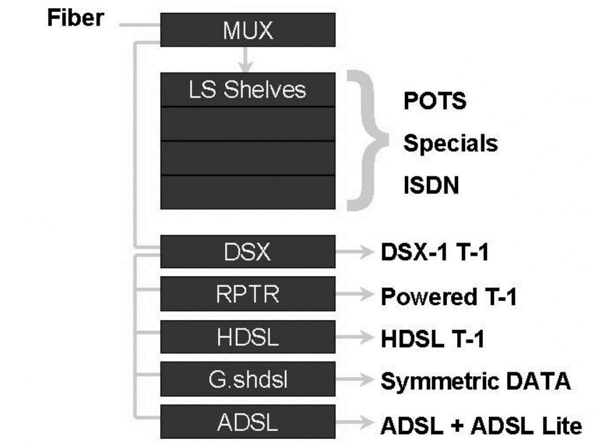

The forecasted development in the current network must be considered in building out the network. A network planner in suburban or rural areas must predict or forecast what will happen in the outlying areas. Typically there is not much time to react. Telephony and datacom demands must be forecast, solutions that meet the forecast must be deployed, and then the next problem must be addressed. This process is expensive, time consuming, and difficult to manage (see *Figure 4*).

For example, in a town there might be three areas. One would have accelerated development, perhaps a business park, subdivisions, and campuses with great opportunities to sell a great deal of ADSL, T1s, and, of course, POTS. Another area would have an existing development with services outside the carrier serving area along a DLC to provide POTS at that site. The third area would be a commercial area with speedy development. Consequently, only T1-type services and POTS sites and services would be needed.

The question then becomes how to serve these different areas. In the traditional model the various products that are needed would be deployed in those areas and the services marketed. For example, a DSLAM, T1s, and DLC would be carried out to the accelerated-development area. In the accelerated-development area with commercial development, a repeater bay with a DLC would be placed to handle all the needed services. The rural subdivision area of existing development would remain unchanged and would be served as is.

FIGURE 4

Forecasted Development for the Box-on-Box Solution

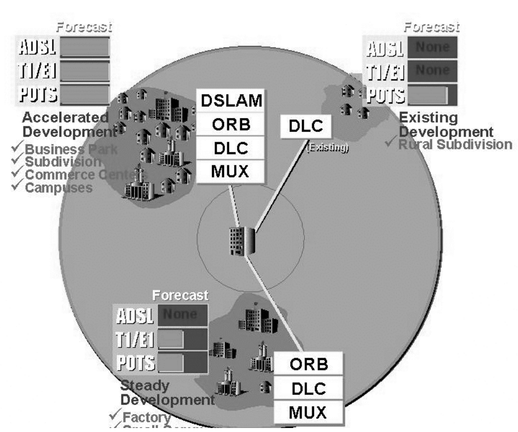

In actuality, many unpredicted developments can occur between the planning, deployment, and marketing cycles. In the accelerated-development area, for instance, cable modems might arrive and undercut the market for the planned services or they may take up some of the demand for the T1 services in that area. In this case, while the POTS is able to serve the basic needs, the DSLAM would remain empty and the repeater bay would be underutilized, resulting in a stranded investment. In the commercial area it may turn out that those businesses prefer ADSL to meet their business needs as opposed to T1s. As a result, despite the T1 equipment placed there, the actual need is for DSL–type services. In the rural subdivision area with the DLC, it might turn out that a business park moves in and demand vast increases for ADSL in the neighborhood, with T1s for local business and POTS. The original planning has left this area under-built and as a result demand goes unmet. Once it is netted out by deploying overlay solutions, it becomes more likely that it will either be over-built or under-built rather than hitting the exact build target. There are real cost and revenue penalties for these misbuilds.

IMAP Advantage

One view is that an IMAP brings a key advantage—integration. The integration lowers the cost of deployment and the risk over the long term. The next-generation DLCs with DSLAM and multiplexer functionality together represent the industry definition of IMAPs. The result is a full-service platform for remote deployments. Whether the transport company has fiber, copper, or radio, the services can be any in and any out, including POTS, specials, T1, ISDN, DS–T1, powered T1, G.shdsl, and ADSL. Plug-in cards enable the addition of whatever services are needed at a given site.

Benefits to the carriers from deploying these kinds of platforms include the following:

- A smaller footprint, because it is all contained in a single product as opposed to multiple different boxes
- Lower power as a result of the ability to integrate all of these solutions
- A single management system to run a full suite of services from a single point, which provides cost savings in the long run
- Loop testing and test access, which will increase the success rate of deployment of DSL types of services
- Integrated voice and data transport that allows the transport part of the network costs to be kept to an absolute minimum

Using broadband integrated multiservice access, the entire serving area can be turned broadband-ready cost-effectively with an IMAP. This can solve the problems for the entire geographical area. By providing plug-and-play, on-demand broadband at any location, IMAP eliminates the digital divide.

In the situational model described above with three development areas, the forecasted deployment now, rather than box-on-box solutions, would use IMAPs to serve the entire area. As a result the actual growth that differs from the forecasted growth is not a problem. It is merely a matter of plugging the appropriate cards into the system to allow it to deliver the necessary services at the necessary time because the carrier deployed a fully integrated broadband multiservice access solution. No over- or under-building takes place. In addition, the system can be managed fully from a single element-management system. All the problems are solved by having a platform that only needs a plug-in module to provide any of the needed services.

The box-on-box solution appears attractive when there are many ports optimized for the maximum price per port. In reality, however, the network often becomes very over-built and underutilized, and there is a high, erratic cost of ownership. In fact it may be nearly impossible to obtain recovery costs at the low line sizes found in remote terminals. This is the step function wherein every addition of a service requires deployment of the additional box to support those services (see *Figure 5*).

Profitability becomes nearly impossible. In areas of accelerated development, there will be demand for increased services with respect to time. If the carrier under-built the network, it will have to add hardware, software, management, logistics, and training incrementally to meet the demand and deliver the services. The carrier, having waited until demand appears, is not in a favorably competitive position against competitive local-exchange carriers (CLEC) and cable companies. It cannot respond quickly, and it loses revenue and market share. With an IMAP, though, there is low scalable cost of ownership. Adding new services only requires the appropriate line cards, and the costs incurred are the services deployed. This allows the continued profitable growth of a full suite of services no matter what the take rates are and what the mix is at any given site (see *Figure 6*).

Box-on-box solutions work with large numbers where future demand is certain to ensure the anticipated return on investment. IMAP is optimal for remote and/or small sites with uncertain future demand where the mixes may vary dramatically. Consequently, deploying IMAP solutions for service and cost-optimized network over the long term is preferable.

FIGURE 5

High Cost of Box-on-Box Solution

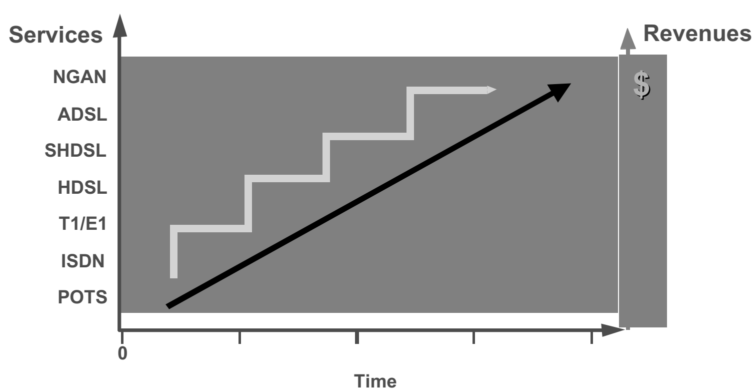

FIGURE 6

The IMAP Solution and Profitable Growth

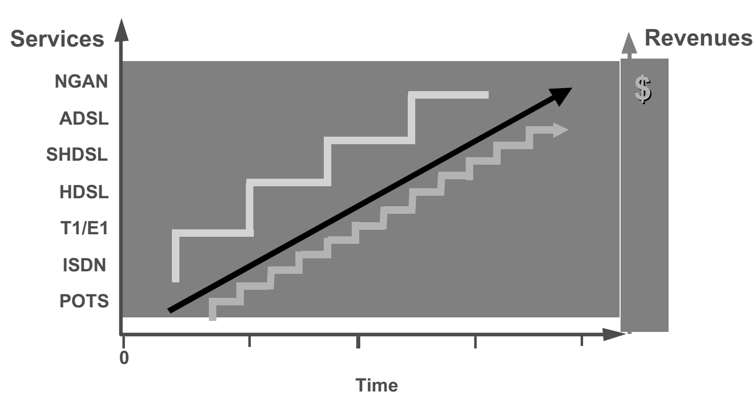

Scaling DSL Profitably

Eric Andrews
Vice President, Marketing
Turnstone Systems

The loop-management business helps service providers more efficiently deploy and maintain digital subscriber line (DSL) services. The focus is on profitability and reducing the costs, particularly in the upfront processes. This paper discusses issues and challenges for DSL service providers.

DSL Service-Provider Challenges

Several issues frequently arise for DSL service providers. One is profitability and the actual cost to get a customer up and running, as high as $900. At a rate of $40 per month in revenue, it will be difficult to make a profit on that customer. There is a concern for profitability and how quickly profitability can be reached with a service offering.

The second issue is coverage. Many service providers are interested in trying to have a pervasive service. It is very frustrating for them to deny service to a customer because they cannot reach that customer because DSL is not available through the digital loop carrier next to that customer's home. It is even more frustrating when customers are told service is possible, and then service in fact cannot be provided, or vice versa. In this respect accuracy of information up front and coverage are serious issues.

A third issue is volume. DSL is an emerging market, still in the early stages. But many people are interested in DSL becoming a mass-market phenomenon; making that migration from a specialty service to a mass market will require significant maturing on many of the back-end operations support systems (OSS) and automated-provisioning systems. Provisioning problems currently hamper moving volume up.

Operational Issues for DSL

Many operational challenges exist for DSL service providers. It is not as simple as turning on a voice switch port and then, suddenly, there are plain old telephone service (POTS) on that line. Many different technologies are involved. Multiple providers' infrastructure, data, Internet, and POTS as well as substantial network activity—loop, splitters, and customer-premises equipment (CPE); digital subscriber line access multiplexer (DSLAM) and asynchronous transfer mode (ATM) backbone; and routers and concentrators—impact operations processes in terms of service negotiation, activation, and assurance. ATM switches, DSLAMs, Internet protocol (IP) configurations, and routers must all be configured. There are outside-plant, copper technology, and cabling issues. Many different things are rolled into a DSL service, and this complexity leads to operational challenges.

Most of the regional Bell operating companies (RBOC) in the United States are currently deploying DSL as a special. Everything that is not POTS tends to be specials, which is typically a high-touch, high-volume, or high-revenue type of application (such as T1 lines).

The challenge is that DSL is, in fact, not high revenue. Marketing DSL to the residential community as a commodity service means revenues of perhaps $40 per month, as opposed to hundreds of dollars per month on that line with another special. As a result there is a need to migrate DSL from the specials mentality to more of the turnkey, flow-through provisioning systems such as a POTS–type service that is easy to deploy, troubleshoot, and so on. This is a significant migration.

Applying Loop Management to DSL

Loop management helps streamline provisioning in the deployment of DSL, and thus should be applied as early in the process as possible. Viewing DSL testing as a troubleshooting tool is not optimal. By the time a customer reports a problem in the network, loss has already occurred because money is wasted trying to find the problem and troubleshoot it. The key is to apply many of these techniques early in the process—during the ordering and provisioning processes—to eliminate troubleshooting calls. By the time a customer receives the service, he or she should not be calling with other questions that need to be resolved.

Loop management can fit into each phase of deploying DSL and help reduce mean time to service, increase customer satisfaction, increase revenue, and decrease costs. At the ordering phase, prequalification can be used; at provisioning, automated-service verification can be used; and at the operations phase, troubleshooting and maintenance can be used. In the ordering phase, a customer calls to order DSL from a service provider. The first question is whether service can or cannot reach that customer. That answer is not always accurate. There are many challenges in trying to get accurate information up front regarding whether or not that customer can actually get service. A technique called prequalification can assist in that process.

One of the problems is that although there are many inventory records on what is available in terms of the copper plant, these records were designed around telephony service and POTS, and they do not translate well to DSL. Those old records and databases are often not useful when used as a source of information for DSL customers—the line may turn out to have a loading coil on it, or it may turn out to be too long, or some other problem may occur. The current loop records are not very accurate, and techniques such as mapping software are being used. Legacy test systems, however, are not effective. Mapping software is limited and does not reflect actual loop path, while handhelds require significant manual labor. Traditional POTS testers are not designed for DSL testing. The result is false positives and false negatives—approximately 30 percent of negatives are false and 20 percent of positives are false. False positives cause wasted effort trying to deploy service to a customer who cannot get it, and false negatives lead to missed revenue opportunities. Thus prequalification becomes a significant way to recapture money being left on the table. Accurate prequalification improves operational efficiency through target marketing to a prospective customer base and avoidance of cost and frustration related to false positives and negatives.

In the prequalification solution, a DSL loop-management system is connected to existing voice switches to test all phone lines and identify those that are eligible for DSL in terms of loop length, load coils, alternating current (AC)/direct current (DC) characteristics, and narrowband noise. Creating this software database provides an accurate picture of the copper plant from a DSL perspective—how many of these lines can actually support DSL and at what rates. Eligible candidates for DSL can be predetermined. This avoids the uncertainty of labor-intensive handhelds and works with back-office components to maintain accurate loop records. When a customer calls, the database can be accessed with a Web interface, for instance, from the order-management system to determine exactly whether or not service is possible (see *Figure 1*).

After prequalification on the ordering phase, the next step is provisioning. The loop-management component fits into that process through automated service verification. The self-install model can help in this area, but other steps are needed to help streamline the entire provisioning process. Potential problems include modem trains without provisioning for the end-to-end ATM virtual circuit; misconnection at the main distribution frame (MDF), perhaps with tip and ring crossed on a pair; CPE configured for virtual path identifier (VPI)/virtual channel identifier (VCI) zero and 35, but a DSLAM configured for eight and 35; cabling errors in the central office (CO), perhaps with the POTS line connected to the wrong splitter port; loop records showing a customer's loop should support asymmetric digital subscriber line (ADSL), but in reality the loop is too long; and router transmission control protocol (TCP)/IP configurations that do not match the modem's configuration, making IP connectivity impossible. As a consequence, the end-user is forced to help with the troubleshooting. Human processes lead to human errors. Automated tools are needed to verify service and streamline provisioning.

Proactive management is required to eliminate problems through service verification before a customer receives a modem. Customer experience can be simulated before shipping a modem by verifying that cabling in the CO is correct, qualifying copper loop, and verifying upper-layer network configuration. If a problem is found at this early stage it can be handled cost-effectively. Results are reduced delays, reduced costs, and more satisfied customers.

Typically the problem is coordination because many different groups are involved in provisioning DSL. Some are setting up the ATM circuits, some the DSLAMs, and some the outside plant. All these organizations must time everything correctly to synchronize on the day the modem is delivered to the customer. In the provisioning and service verification, the final step before shipping a modem to the customer, then, should be to confirm not just the physical outside plant but also all the different protocol configurations (IP, ATM, protocol encapsulations, and so on). Cabling should be verified by checking the DSLAM signature to ensure proper cabling, checking for improper termination at MDF, and verifying that the phone number

FIGURE 1

Prequalification Solution

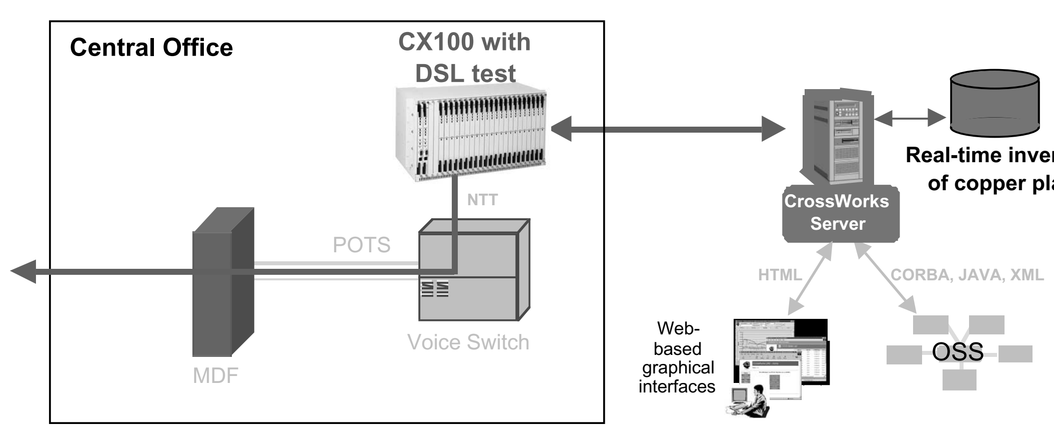

FIGURE 2

Cable Verification and Loop Qualification

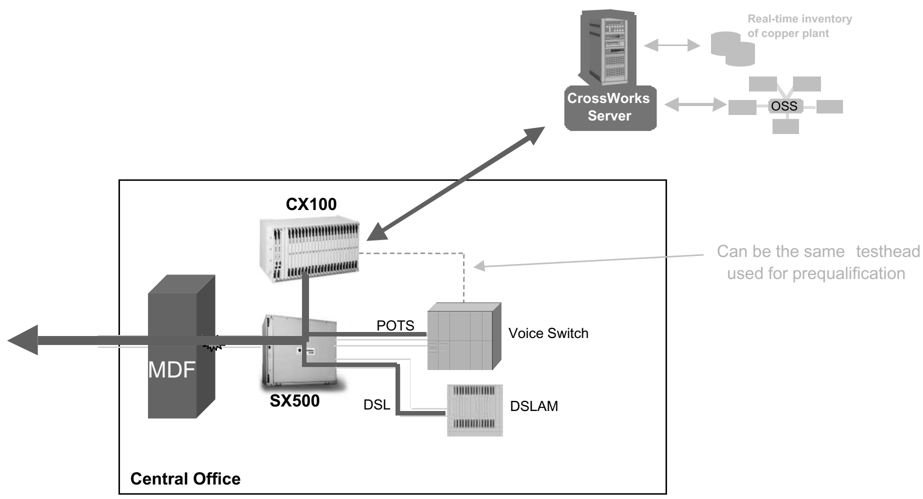

matches the DSLAM port—automatic number identification (ANI). Loop qualification should include testing the loop (bridge taps, load coils, loop length, and so on) and comparing the test results to benchmarked information. The end result will be more satisfied customers and faster troubleshooting (see *Figure 2*).

Beyond just the physical plant, service verification can confirm the upper-layer protocols. Self-install can be simulated prior to shipping CPE to a customer through verification of the IP layer (IP ping, trace route, and address resolution); verification of encapsulation (for example, point-to-point protocol over Ethernet [PPPoE], point-to-point protocol over asynchronous transfer mode [PPPoA], and so on); verification of ATM signaling channel F4/F5 operations, administration, and maintenance (OAM), and ATM ping; and verification that the DSLAM line card trains (see *Figure 3*).

If a problem is found, it can be addressed instantly. This eliminates frustration for the customer and also saves money for the service provider that can effectively deal with these issues up front and ensure that self-installs are accurate. Those two steps are very significant in terms of reducing the cost and complexity for DSL deployment and installation.

The final step is ongoing operations and maintenance. Historically, any service ever deployed to a mass market has included remote-testing capability; indeed, it is more expensive in the long view not to include testing. For instance, the POTS network, the ultimate commodity network, has testing capability and automated systems.

FIGURE 3

Service Verification

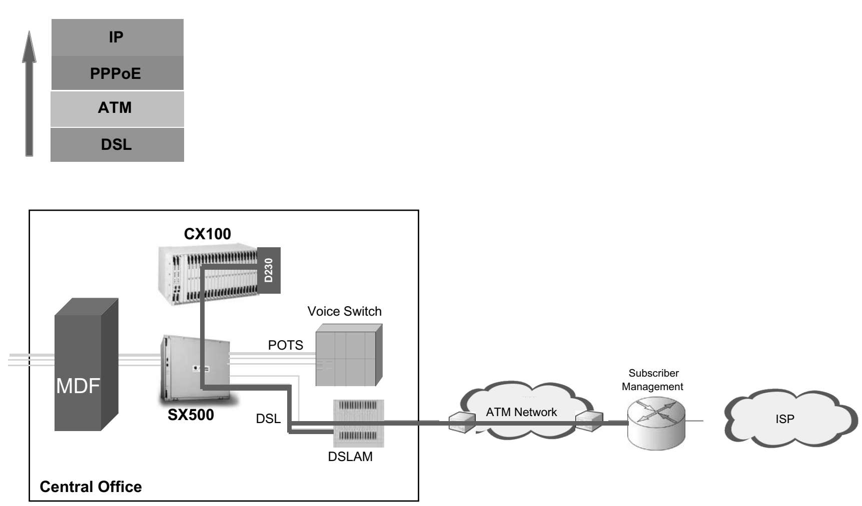

Dispatching a technician every time there is a trouble call is extremely expensive. That is the case with DSL currently, and that needs to change. The migration to a mass-market type of service requires automated testing capability built into the infrastructure.

One of the issues is single-ended testing. This helps avoid truck rolls unless absolutely necessary. Another issue is accurate fault isolation to identify where the problem is—an outside-plant issue, an ATM issue, or an IP–layer issue—as opposed to exactly what the problem is. The appropriate groups then can be involved in solving the problem. Finally, testing also must work with splitters. As testing infrastructure begins to be deployed into DSL networks, problems are emerging with testing through the splitter. There is limited testing from the equipment side of the splitter. The DSL side cannot perform such basic loop-qualification tests as loop length. The POTS side cannot perform wideband tests; POTS test systems, for instance, are not designed to handle DSL–specific issues such as load coils. The same is true with testing through DSLAM (see *Figure 4*).

One key issue is test access at the right places in the network. The ability to identify network and loop issues requires visibility. A concept called a smart splitter provides metallic access on all sides of the splitter. This eliminates the problem of the splitter impeding ability to view what is occurring on the circuits. Clear visibility is necessary, both out toward the loop and back toward the network, to do an effective job of loop management in DSL networks. A smart splitter offers cost-effective integration of metallic access and splitters, simplified cabling, and full loop-management capabilities (see *Figure 5*).

DSL and the Remote Terminal

The remote-terminal (RT) space is a challenge for DSL service providers. Many people behind digital loop carriers currently cannot get DSL, even if there is a DSLAM in the CO. Those digital loop carriers are fiber-fed from the CO, and thus there is no copper going back to that CO. Service providers are very interested in reaching 100 percent of the potential customer base, and to do that they must find ways to get DSL into the RT space. DSL deployments, in fact, are emerging in the RT space. The same loop-management issues apply—prequalification, automated service verification, and troubleshooting and maintenance. In the RT space, however, everything is magnified because of the higher number of locations. There might be 10,000 RTs, rather than 1,000 COs. In addition, truck rolls are extremely expensive, and visiting the RT to do something is costly. As much automation as possible is necessary.

In the DSLAM approach to DSL in the RT space, service providers place a DSLAM–type device, basically a separate cabinet, side by side with the existing cabinets. That device is DSL focused, and at the same time it is focused on all the data services. The splitters, the testing, and the DSLAM all are put in that cabinet, and it is cabled across to the existing digital loop carriers for POTS. This is seen in certain RBOCs and incumbent local-exchange carriers (ILEC).

A different approach is to use an integrated, next-generation digital loop carrier. Some digital loop carriers have DSL line cards that can be plugged in, and some new products are emerging that integrate POTS and DSL on their report. Many developments are underway. In either scenario the same concerns remain—visibility to test those loops, both

FIGURE 4

Limitations of "Dumb" Splitters

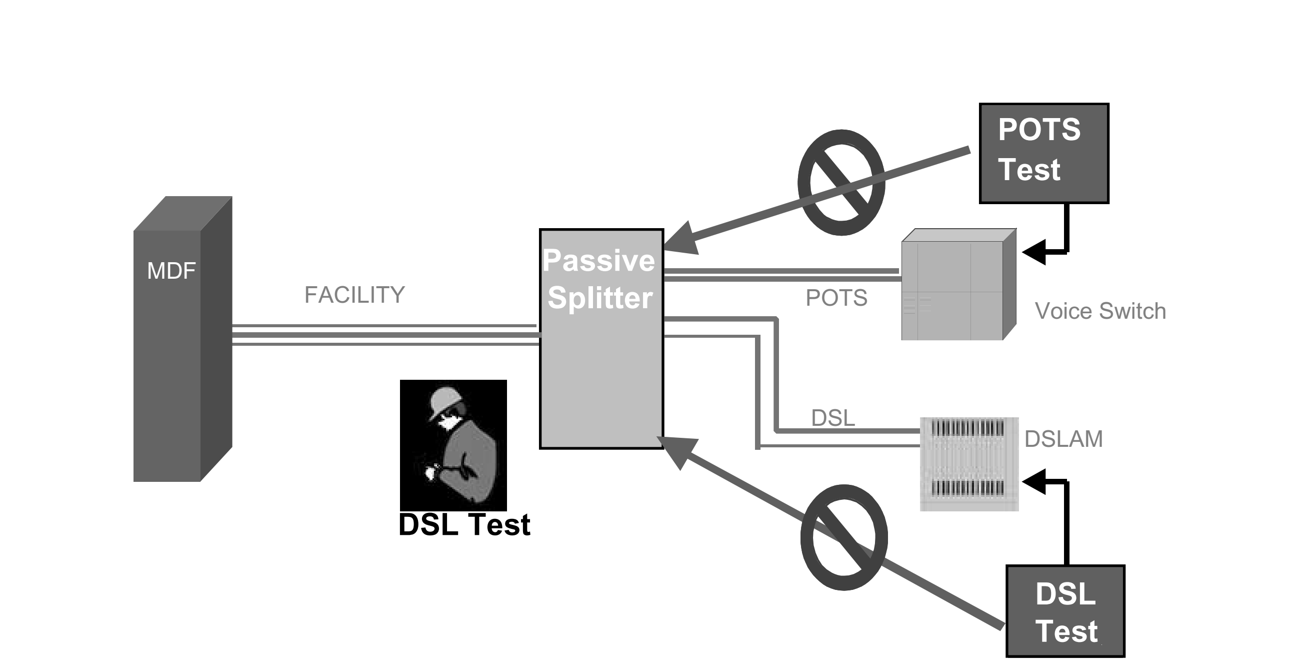

FIGURE 5

Smart Splitter

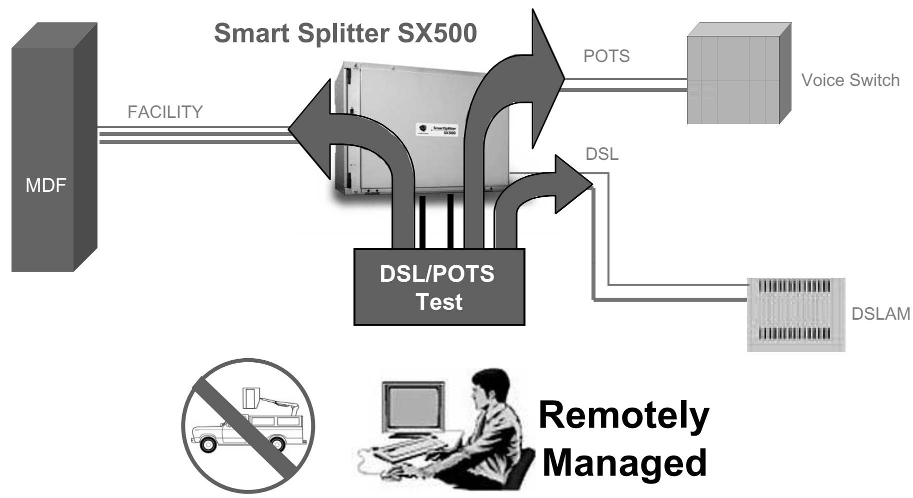

out toward the outside plant as well as back toward the network, and the correct infrastructure to perform the necessary testing and automate it.

Conclusion

Loop management is critical for profitable DSL. DSL can be deployed without loop management, but it will cost more in terms of truck rolls and other expenses. With loop management the service providers can accelerate revenue growth through fewer delays and errors in deploying DSL, lower operation expenses by reduced requirements for troubleshooting and on-site labor, and improve customer satisfaction by increased first-pass success and shorter time to deploy and repair.

The Architectural and Economic Advantages of 1D MEMS–Based Wavelength Cross-Connects

Douglas J. Arent
Director, Strategy and Business Development
Network Photonics

Bevan Staple
MEMS Development Engineer
Network Photonics

Multiplexed domain switches enable low-cost, simple, reliable, and scalable architectures for a wide variety of wavelength selective cross-connects, dynamically reconfigurable add/drop multiplexers (ADMs), and hybrid optical core switches (OCSs).

As dense wavelength division multiplexing (DWDM) becomes ubiquitous in long-haul and metro core networks, carriers are looking to optical switching to manage the ever-growing number of wavelengths and to automate wavelength routing, restoration, and grooming. To meet these requirements, OCSs need to scale to accommodate high wavelength counts, high tine division multiplexing (TDM) bit rates and various types of transport protocols, such as synchronous optical network (SONET) and Gigabit Ethernet (GbE).

O–E–O versus O–O–O Architectures

OCSs, also known as optical cross-connects (OXCs), can generally be classified into three categories: 1) opaque, electronic core switch, 2) opaque photonic core switch, and 3) transparent photonic core switch. Opaque, optical-to-electrical-to-optical (O–E–O) OCSs offer the ability to groom services down to subwavelength bit rates such as synchronous transport signal (STS)–1. This grooming capability results in scalability challenges due to bit-rate limitations and high power consumption of the electronic switching fabric. Furthermore, O–E–O OCSs are bit-rate and protocol dependent (opaque), which limits flexibility and requires forklift upgrades.

Opaque photonic switches combine O–E–O conversions and optical switch fabrics, addressing the power and scalability issues of O–E–O OCSs while remaining bit-rate and protocol dependent. In contrast, transparent or optical-to-optical-to-optical (O–O–O) OCSs, where no O–E–O conversion takes place, are bit-rate and protocol transparent. O–O–O switches thus offer more flexibility to carry diverse protocols and bit rates and are particularly useful in topologies where TDM rates may be upgraded without requiring O–O–O OCS upgrades. Furthermore, since OOO switches route signals using passive optics (without O–E–O conversions), their power consumption can be significantly less than that of equivalent scale O–E–O switches.

O–E–O and O–O–O OCSs do not necessarily compete with each other. When viewed from a network-wide perspective, O–E–O and O–O–O switches offer complementary benefits. O–E–O switches are best used at the transport edge, where the grooming of services is needed, and O–O–O switches are best used in the core of the network, where management at the wavelength level is needed.

OOO Switch Technologies

While multiple technologies are being developed for the core of O–O–O switches (e.g., including bubbles, MEMS, nonlinear optical materials such as Lithium Niobate, holograms, and liquid crystals), the micro-electromechanical system (MEMS) has emerged as the leading technology of choice. Generally, these MEMS switches are classified as two-dimensional (2D) or three-dimensional (3D). The 2D planar arrays, also known as N^2 architectures (since N^2 mirrors are required to switch N wavelengths) are now available in low port counts (e.g., 2x2, 4x4, and 8x8), but the 2D architecture has limited scalability (e.g., 32x32) due to diffraction challenges and limitations on the size of the MEMS die. 3D MEMS architectures are currently under development in an attempt to achieve high port counts without sacrificing optical performance. However, 3D architectures, also known as 2N architectures (since 2N mirrors are required to switch N wavelengths) pose complex and expensive control and packaging challenges.

In 3D architectures, since each micromirror must have multiple stable positions, a sophisticated closed-loop servo system with feedback must be implemented for each mirror to ensure that the mirrors move to and remain in the desired position once commanded. Using a closed-loop control scheme implies that one must monitor the beam positions over time, space, and environments and actively feed this information back to the system in a timely fashion. The required 2N servo systems force solutions that are large in size and power consumption, and have high costs.

WXC Architectural Advantages

In contrast to OCSs, wavelength selective cross-connects (WXCs) are transparent, optical multiplex-domain core switches that offer significant advantages over OCS–based architectures.

WXCs offer multiple advantages over OCSs. As shown in *Figure 1*, a WXC natively switches wavelengths in the multiplexed domain. The multiplexed DWDM wavelengths enter the WXC and are switched among input and output fiber ports, and then exit the WXC as multiplexed signals on the desired fiber. Switching in the multiplexed domain eliminates the need for external demultiplexing and multiplexing, as well as multiple O–E–O conversions. The WXC offers significant fiber management and cost savings by requiring only a *single port per DWDM fiber*, as opposed to a single port per wavelength in OCS–based systems. For example, each fiber in the 3x3 WXC shown in *Figure 1* may carry 80 wavelengths, thus the 3x3x80 architecture offers effectively 240x240 wavelength switching capacity, but only requires six fiber ports.

FIGURE 1

An All-Optical WXC

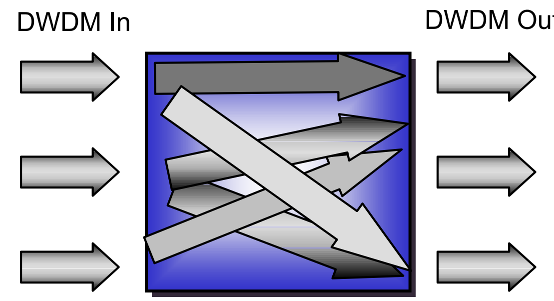

As a purely photonic cross-connect switching in the multiplexed domain, a WXC has extremely small power requirements (< 10 watts), in contrast to many kilowatts for traditional O–E–O OCSs. By eliminating the external demultiplexing/multiplexing and O–E–O conversions and reducing fiber-management complexity, the WXC has an extremely small footprint. Because it photonically switches optical signals, the WXC is both bit-rate and protocol transparent. Wavelength bit rates may be upgraded at the transport edge (i.e., from 155 Mbps to 40 Gbps), while the WXC does not require any additional operational input or change.

As summarized in *Table 1* for comparable OCS and WXC configurations, significant cost savings complement all of the noted advantages. Capital expenses are reduced through the elimination of O–E–O conversions, closed-loop servo controls and the required demultiplexing/multiplexing equipment. Operational expenses are reduced through lower space, power, and management costs. The dynamic nature of the WXC allows for remote software control and advanced protection features, offering significant advantages from an operational level and in the ability to rapidly provision services. The CrossWave™ technology enables all of these advantages.

CrossWave™ 1D MEMS Technology

CrossWave integrates demultiplexing/multiplexing and switching into a single operation, in contrast to other all-optical switch-based systems that require external demultiplexing/multiplexing. This powerful functionality is achieved through a hybrid integration of one-dimensional (1D) MEMS and a wavelength dispersive element. Within CrossWave, wavelengths from input fibers are spatially separated by the dispersive element. Each separated wavelength is focused onto the surface of a micromirror in the 1D MEMS array where the switching occurs. The independently switched wavelengths are then recombined into DWDM signals and exit into the output fibers.

TABLE 1

Comparison of Key Features of a 512x512 OCS and a 8x8x64 1D MEMS–Based WXC

	OCS	**WXC**
# of DWDM fibers switched	8x8	8x8
# of λ's per fiber	64	64
Switch configuration	512x512	8x8x64 (512x512)
# of physical ports	1024	16
# of DWDM mux/demux channels, transponders, and receivers	1024	0
Equipment footprint		
Power required		
Total cost	**$$$$$$$**	**$**

The CrossWave 1D MEMS technology consists of a single, linear array of micromirrors. The 1D architecture is highly scalable in that only one mirror is required to switch each wavelength (e.g., an N-type switch versus N^2 for 2D or 2N for 3D architectures). Furthermore, each mirror is positioned accurately in one of several highly stable positions within switching times of less than a millisecond. The 1D linear array of micromirrors is powered with simple digital electronics in an open-loop configuration. With digital positioning, there are no complex servo systems necessary to ensure stable and repeatable positioning of the micromirrors. Furthermore, digital control eliminates the need for complex sensing and low-noise electronics. This results in a dramatic reduction in size, cost, and power consumption compared to 2D and 3D systems, while increasing system reliability. An example of the 1D MEMS architecture is shown in *Figure 2a*, in which N mirrors address N optical signals.

Another key difference between the 2D or 3D and 1D MEMS switches is the mirror packing density of the die. While 2D or 3D MEMS typically occupy much of the real estate on large silicon die, 1D MEMS utilize much smaller mirrors arranged in a linear configuration, occupying only a small fraction of the die. This results in higher manufacturing yields due to lower susceptibility to contamination and handling damage. In fact, the die dimension and layout for 1D MEMS are driven by packaging needs instead of the mirrors, thereby increasing the yield and reliability of the overall packaged device.

Summarized in *Table 2*, the key features of MEMS–based optical components indicate the relative advantages and disadvantages among 1D, 2D, and 3D devices. 3D MEMS devices, while holding strong promise for scalability, offer a number of challenging issues such as mirror control, packaging complexity, and overall reliability. 2D MEMS devices have been deployed in small cross-connects but do not scale well and also offer packaging challenges. The combination of Network Photonics' CrossWave 1D MEMS architecture, ease of manufacturing, ease of control, simple packaging, operating robustness, and scalability offers an elegant, reliable, and low-cost solution.

Physically, CrossWave is fully athermal, eliminating any requirement for active temperature control and extremely low power (e.g., < 1 watt per device). Network Photonics leverages Crosswave technology to produce highly reliable, low-cost optical switching and transport equipment to meet the stringent requirements of the telecommunications industry.

Network Applications

CrossWave-based WXCs can address a number of applications in a variety of network topologies, ranging from linear to ring to mesh topologies. In ring topologies, distributed or centralized WXCs or ring interconnects enable wavelength routing, protection, and restoration while eliminating many O–E–O conversions. Similarly, in mesh networks, WXCs can be applied to all-optical wavelength routing. In linear net-

FIGURE 2a

An Overview of Network Photonics' Linear 1D MEMS Array

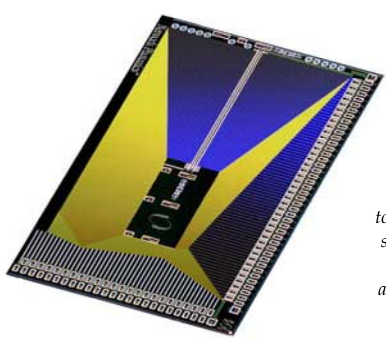

The vertical line in the top/center of die shows the single row of 1D MEMS mirrors. The squares around the edge of the die are bonding pads.

FIGURE 2b

A Scanning Electron Micrograph (SEM) of a Section of the CrossWave™ 1D MEMS Array

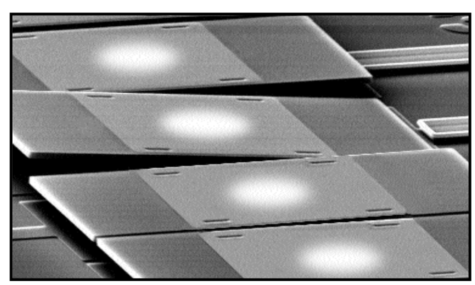

TABLE 2

Comparison of 1D, 2D, and 3D MEMS Architectures

Technology/Feature	Crosswave™ 1D	2D	3D
Mirror Control	Simple	Simple	Difficult
Ease of Manufacture	Easy	Complex	Very Complex
Packaging Complexity	Simple	Somewhat Complex	Very Complex
Reliability	High	High	Low
Scalability, Mirrors per λ	N	N^2	2N
Cost	Low	Medium	High

works, WXCs effectively address two pressing applications and are subsequently described in more detail: flexible add/drop capability in long-haul (LH) and ultra-long-haul (ULH) DWDM networks and hybrid OCS switching.

Backbone Network Applications

One critical challenge in deploying LH and ULH DWDM networks is cost-effectively serving the numerous add/drop sites to the backbone. Since current generation LH and ULH DWDM network equipment allows all-optical transport up to many thousands of kilometers without regeneration, carriers must address the need to add, drop, or groom wavelengths at network nodes in between the end terminals.

Figure 3 depicts a long-reach DWDM network link including intermediate add/drop sites. At the intermediate sites where wavelengths need to be flexibly added, dropped, or groomed for appropriate services, the key challenge is to effectively minimize any perturbation to desired express traffic and cost-effectively add, drop, or groom the required wavelengths. While the required wavelength-level functionality can be addressed with a variety of solutions—including fixed filters, power splitting combined with full demultiplexing and selective extinction, or banded filter architectures—these solutions are neither cost-effective nor sufficiently flexible to address traffic patterns that are neither static nor reliably predictable. Furthermore, cascading multiple elements in a single link requires excellent optical performance, such as flat-top passband and dispersion free design. Ideally, service providers prefer to deploy a flexible add/drop network element to effectively address low initial cost requirements, low operating expenses, required flexibility, and scalability to handle changing and unpredictable traffic demands.

Dynamic Wavelength Selective Add/Drop Cross-Connects

At the intermediate nodes, the wavelength selective add/drop cross-connect allows one to individually address any wavelength or set of wavelengths within the multiplexed cross section, thus allowing dynamic, flexible reconfiguration of express and add/drop traffic. In these configurations, wavelengths can be reassigned from the express path to the add/drop paths with no effect on the wavelengths remaining in the express path.

Functionally, multiplexed DWDM signals enter, are switched among fiber ports, and then exit the add/drop WXC as multiplexed signals on the desired fiber. For express traffic, switching in the multiplexed domain eliminates the need for external demultiplexing and multiplexing, and multiple O–E–O conversions. This capability directly eliminates 50 to 70 percent of the costs of intermediate add/drop nodes compared to alternative solutions.

For drop traffic, small wavelength count demultiplexers may be used to isolate individual wavelengths from the multiplexed signal. In addition to the advantages previously outlined, including bit-rate and protocol transparency, extremely low power, and a remarkably small footprint, the add/drop WXC offers significant fiber management and cost savings by requiring only a *single port per fiber*, where each fiber can carry a large number of wavelengths.

Hybrid OCS

Similar to the wavelength selective add/drop cross-connect, hybrid OCS architectures allow dynamic, flexible reconfiguration of express or grooming traffic. In these configurations, wavelengths can be reassigned from the express path to the grooming switch fabric with no effect on the wavelengths remaining in the express path. By selectively grooming only the required services, the hybrid OCS offers effective wavelength management as well as cost-effective and appropriately sized grooming capabilities, resulting in lower overall costs, space utilization, and power consumption.

CrossWave Technology

CrossWave multiplexed switching enables a new paradigm in all-optical networking, making fully functional, highly

FIGURE 3

Extended-Reach LH or ULH DWDM Network with Dynamically Configurable Intermediate Add/Drop and Grooming Nodes

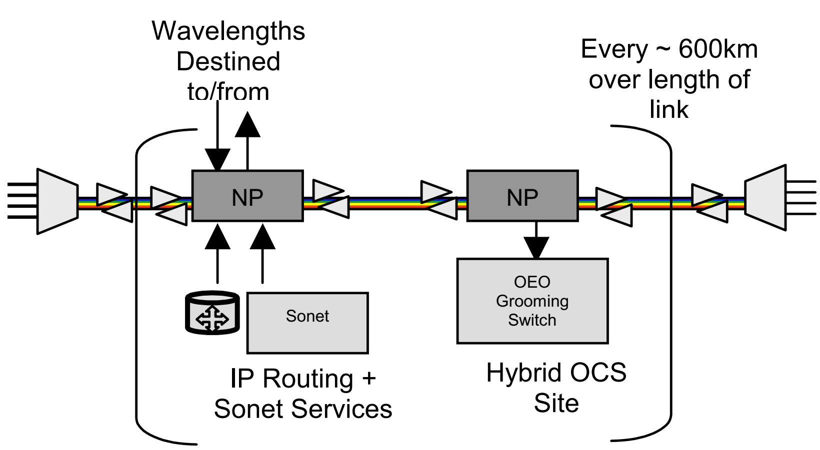

The left-hand node offers dynamic wavelength routing to/from other service nodes as well as SONET and data services. The right-hand node offers dynamically configurable hybrid OCS wavelength grooming services. The intermediate nodes are typically spaced ~600 km to serve lower-tier cities and mid-span connection points.

reliable multiplexed switching possible while enabling new economic efficiencies required by service providers. Utilizing a highly integrated architecture of 1D MEMS and passive optics, CrossWave leverages advances in high-reliability volume manufacturing, 1D MEMS scalability and reliability, fast switching times, and packaging simplicity to bring new features to all-optical DWDM networks. CrossWave multiplexed domain switches enable low-cost, simple, reliable, and scalable architectures for a wide variety of scalable WCXs, dynamically reconfigurable ADMs, and hybrid OCSs. In long-reach DWDM networks, flexible add/drop network elements and hybrid OCSs may be deployed to effectively address low initial cost requirements and low operating expenses as well as to provide the required flexibility and scalability to handle changing and unpredictable traffic demands.

Intelligent Agents on the Web: An Overview of Design and Applications

Sunil Babu
Graduate Student, Computer Science, College of Engineering and Sciences
Louisiana Tech University

Vir V. Phoha
Associate Professor, Computer Science, College of Engineering and Sciences
Louisiana Tech University

Abstract

Since its inception, the Web has become more popular with each passing day. New concepts are being researched and implemented for efficient usage of the Web. Web intelligence is one new, emerging area, which promises to bring together artificial intelligence (AI) technology and advanced information technology (IT) for the next generation of Web-empowered products, systems, and activities. Intelligent agents that use Web intelligence are software tools with AI techniques whose application brings intelligent features to the Web and Internet. This paper introduces intelligent agents (IAs), their design, and some of the applications in which they are being used.

1. Introduction

IAs are autonomous and adaptive software programs that accomplish their goals by proactively executing predetermined commands. Agents are software tools that act to help users to perform certain task efficiently. In a broad sense, an IA can be thought to be a back-seat driver that gives suggestions at every turn or as a taxi-driver who transports a passenger to his destination. There are advantages in preferring the usage of agents [8]. To name a few, agents adapt dynamically to their environment, execute autonomously, are robust and fault-tolerant, and encapsulate protocols. Agents are used in e-commerce where they can act on behalf of customers; they are used as personal assistants to assist users; they are used in secure banking, in distributed information retrieval, in telecommunication networks services, and in parallel processing; they are used for information dissemination, such as automatic news and software updates for vendors; they are used to monitor a particular information source, and they are used in workflow management, in air-traffic control, in network management, and in smart databases-to name a few applications.

Etzioni and others [13] have reasoned that since the development of analyzing tools have not kept pace with the information growth, the information superhighway has not been forgotten by people and still seems to be prominent in human lives. The growth of agents pave the way for this information superhighway disappearance, as they have the property of abstraction (the technology and resource agents use are abstracted from the user) and distraction (the agents' inherent abilities distract a person from performing complex and tedious tasks).

2. Intelligent Agents

Wooldridge and Jennings [1] define an agent as any software system, possibly including a hardware system, with the following properties:

- *Autonomy:* Agents operate on their own without the intervention of any humans. They have control over their actions and their internal states.
- *Social Ability:* Agents interact with each other exchanging useful information in a pre-defined language (agent communication language) of their own.
- *Reactivity:* Agents gather information from their environment and respond appropriately to the changes that occur in the agent environment.
- *Proactive:* They are proactive, performing their desired goal with their own initiative.
- *Mobility:* Depending on their application, agents are mobile, moving around the electronic world and exchanging information with their peers.
- *Veracity:* Agents do not communicate any false information that is not supposed to be communicated by them.
- *Benevolence:* Agents do not have conflicting goals, and all agents collectively co-operate with each other to achieve the desired goals.
- *Rational:* Agents will act toward achieving their goals and do not act in a manner that prevents them from achieving their goals.

Every agent has the first four properties along with a subset of the remaining properties. A simple agent can be thought of as a program (e.g., a UNIX process) that has a determined goal and exhibits the properties of an agent that were listed.

3. Building an Agent

Wooldridge and Jennings [1] have given steps to design and realize agents. Agent building can be categorized into three main phases.

1. *Agent Theories:* Deals with the specification and conceptualization of agents. It identifies what properties an agent must have and how these properties are formally represented. Agent theories define formal methods for representing the properties of an agent, including methods to represent dynamic aspects of agents. Agent theories must include representational methods for representing the relation between the attributes of an agent as well as defining the cognitive state change in an agent, the effect of the environment on the agent, and the way in which agent information leads to its action.
2. *Agent Architecture:* Deals with the design and implementation of an agent in a computer system. It specifies the decomposition of the agent design into component modules and has a provision for interaction between these modules.
3. *Agent Language:* Deals with the communicational aspects within peer agents. The agent communicating language (ACL) must be capable of communicating the state, properties, goal, information, and behavioral aspects of agents. A few examples of agent languages include concurrent object language, agent-oriented programming language (AOP), telescript, planning communicating language (PLACA), concurrent MetateM, APRIL, the Foundation for Intelligent Physical Agents' agent communication language called (FIPA), FIPA ACL, and MAIL.

3.1 Layered Architecture of Agents

In designing agent architecture, Kendall and others [2] give a layered architecture design for agents. This design has seven layers (see *Figure 1*). The higher-level layers are dependent on lower layers for functionality, and there exists a two-way flow of information between layers. This framework design is reusable, robust, and is based on object-oriented patterns. This framework demonstrates that agents can be designed and implemented using objects. This framework is employed in network-management applications, network-intrusion detection, and enterprise-integration applications.

The seven layers, from the bottom up, are explained in the following:

1. *Sensory:* This layer proactively senses the agent environment and updates the agent information. This layer co-operates with domain-specific interface classes.
2. *Beliefs:* This layer stores the belief of an agent. The input from the sensory layer is used to update the belief of an agent, which will be used for reasoning.
3. *Reasoning:* This layer determines the next step/action of an agent. This layer interprets the beliefs, and, based on the interpretation, an action is initiated or the state of the agent is changed.
4. *Action:* This layer is responsible for executing an action responsible for the goal of an agent. Each action is an intended goal of an agent, which takes the agent toward its final goal.
5. *Collaboration:* This layer is responsible for carrying out interactions and exchanging services and negotiations between agents. It addresses the issue of collaboration and co-ordination of different protocols between agents.
6. *Translation:* This layer is responsible for any translation of different protocols that is being exchanged between the agents.
7. *Mobility:* It supports for the virtual or actual migration of agents between environments. It provides for the cloning of an agent for actual migration into another environment.

4. Taxonomy of Agents

Franklin and Graesser [10] suggest various ways in which an agent can be classified. An agent can be classified based on the subset of the properties listed in Section 2. Every agent has the first four properties. Agents having additional properties are classified differently according to the properties. For example, agents having the property of "mobility" are classified as mobile agents. A hierarchical classification is thus possible based on the property subset. Thus mobile learning agents are of a sub-class of mobile agents.

Other possible schemes of classification are suggested. Based on the task that agents perform, they are classified as information-gathering agents or e-mail-filtering agents. Based on their control mechanism, agents can be classified as rule-based, algorithmic, fuzzy, neural net, or machine learning. Agents can also be classified based on the environment in which they operate, such as software agents and artificial life agents. Agents can even be classified based on the language in which they are written and on the application for which they are written.

Brustoloni and C. Jose [11] provide a three-way classification of agents:

1. *Regulatory Agents:* These agents do not perform any planning. They have knowledge built in them to cope with the stimuli of an environment. They have a set of prescribed actions that satisfies the overall objectives of the agent. An example of this kind of an agent is a thermostat that controls the temperature. These kinds of agents know what to do based on the inputs they receive from their environments. Such agents can be "tunable" according to their environment and "configurable," that is, a pre-determined behavior can be selected.
2. *Planning Agents:* These agents, along with the functionality of regulatory agents, have the capability to plan. Some of the methods used as planning strategies are as follows:
   - Control systems and automata
   - Problem-solving paradigm of AI
   - Case-based paradigm of AI
   - Methods described in operations research
   - Randomizing algorithms

     A common problem in using these approaches is that they are very expensive.
3. *Adaptive Agents:* These agents have the capability of planning along with the ability to acquire knowledge for planning. They are capable of acquiring domain knowledge, which gives the agents the ability to perform new actions. An advantage of this method is the improvement of performance through new techniques. These agents are particularly suitable for an environment that is dynamically changing.

5. Agent Security

In the world of the Internet, agents are used mainly for information gathering or in e-commerce. These agents have

FIGURE 1

Seven-Layered Framework of an Agent

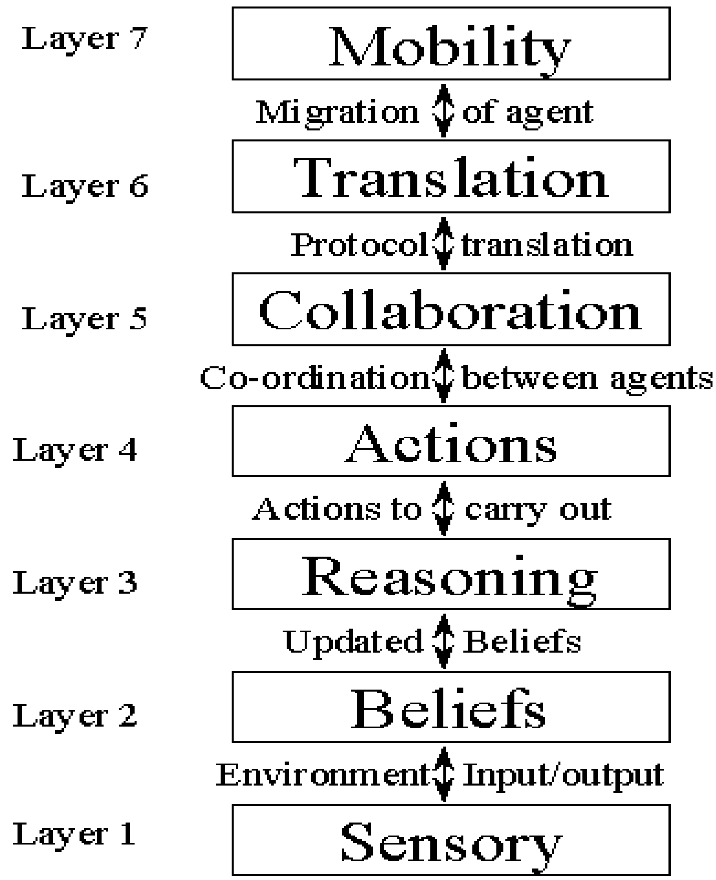

to move about from host to host performing their intended goals. The hosts that they visit may not be trusted and have the malicious intention of attacking the agents. Grimley and others [14] have given ways in which a malicious host may attack an agent and possible solutions for these attacks. The types of security attacks a malicious host may perform on an agent are as follows:

1. *Invasion of Privacy:* The malicious host has access to the agent's memory space and hence can read the potential information present in the agent that was not intended to be shared by the hosts. The data read could be a secret key, electronic money, or the lowest price bid of an item.
2. *Data or State Manipulation:* The malicious host can alter the data or information present in the agent. For example, it could change the cost of an item or, even worse, it could change the agent's belief or commitment so that the agent favors the host.
3. *Code or Control-Flow Manipulation:* The malicious host can alter the code of the agent to suit its needs or desire. Even worse, the host can implant a virus into the code of the agent.
4. *Purposeful Misinterpretation of Agent Code:* Agents are normally coded using an interpretive language such as Java or TCL. The malicious host can change the interpreter itself such that the execution of the code favors the needs of the host.
5. *Cloning:* The malicious host can clone a copy of the agent and analyze its code and behavior. In the next visit of the agent, the host can then perform any of the attacks previously listed.
6. *Denial/Delay of Service or Misdirection:* The malicious host can redirect the agent to a host not in the itinerary of the agent, or the host may not release the agent.

Possible solutions exist to prevent/detect the attacks and involve varying levels of overhead, which determine the compromise balance between the security level and runtime efficiency of the agent. The possible solutions for these attacks are as follows:

1. *Data Encryption:* An obvious solution is to encrypt the agent's data. The agent that does not require its data until it has reached its home encrypts the data and delivers it in the encrypted form to the home base. The home server possesses the decryption key to decrypt the data.
2. *Code Encryption:* In this method, the code itself is encrypted. An encrypted code has to be decrypted on the host before it executes. This poses the same threat as an unencrypted code. An approach to counter this is

called blackbox security, where the execution layer accompanies the encrypted code, and this execution layer executes the agent code on the host machine.

3. *Time-Sensitive Agents:* In this approach, the amount of time an agent executes on a host is determined. If the time is exceeded, then the agent shuts itself or transfers itself to another host. This approach is suitable given the fact that it takes time for a host to examine and evaluate an executing agent for possible attack.
4. *Code Obfuscation:* This technique obfuscates the code so as to make it less readable and comprehendible. This technique removes modularity, splits the meaningful variables, and inserts code that does not alter the goal of the agent. An agent visiting the host again has a different form of code.
5. *Phoning Home:* A method to minimize the damage after tampering is to have the agent phone home and send the data collected. This prevents the loss of data and its disclosure to future hosts. This technique allows the host to determine if the agent is functioning properly, and if not, at which node it has failed.

6. Applications of Agents

Much research is being carried out in applying agents in new areas. Agents are useful tools and are predicted to be ubiquitous in the near future. A common application example of an agent is a personal digital assistant (PDA) [1]. The PDA filters e-mails and sorts them in order of importance. It draws attention to any important message that has been received, like a paper being accepted in a conference. It proactively searches travel agencies and provides a list of economic travel schedules to attend, for example, a conference. All of this is done even before the e-mail is read.

Existing agent applications include finding citations on the Web [3], improving network operation [4], use in electronic commerce [5], Web-browser intelligence [6], and the intelligent profiling of users [7]. In the following sections, we will explore the details of agent implementation for a few important applications.

6.1 Improving Network Operations

A network has to be constantly supervised and managed for any errors or high traffic. Simple network management protocol (SNMP) provides means to poll the network devices to know their performance. This constant polling can lead to more network traffic. To reduce this, intelligent agents are used locally that constantly pool the network device and send a network signal to the manager only if the device is not functioning properly [4]. Agents can even have corrective-action capabilities to rectify any error. Intelligent agents can be used in similar applications, such as managing many branches of a bank. Localized branches have agents present. Agent sends a message to the manager in a higher hierarchy only if any problems, such as a storage error, occur in the localized branch. The manager, which is also an agent, can then dispatch a specialist, in this instance a storage specialist, to rectify the error.

The agents are used in the following network-management areas [4], where the data collected by them is useful in determining the type of the error. The type of data being collected is based on the application of the agent.

- *Performance Management:* Performance monitoring can help managers to understand network behavior. Agents can gather information regarding application usage, client/server performance, and data correlation to determine which of the nodes has peak usage and at what times.
- *Fault Management:* It is important for fault to be identified immediately to reduce the negative impact on the services. The agents can be used to identify the faults, immediately notify a manager, execute automated resolution procedures, and isolate the faulty node for the manager to get at the heart of the problem.
- *Capacity Planning and Reporting:* This service is crucial to deliver service levels to the user. Agents can be used in load balancing, identifying loaded nodes and diverting traffic to lesser-loaded nodes.
- *Security Management:* Intelligent agents help to discover security holes in the system. They monitor the firewall configuration and determine if it is functioning well. They monitor access to secure subnets for security breaches.

6.2 IAs for Electronic Commerce

IA models have been developed to assist in electronic commerce. The multi-agent negotiation testbed (MAGNET) [5] is an e-commerce architecture developed to assist in a business-to-business (B2B) situation where companies outsource production to contractors. Agents have been developed to automatically negotiate and monitor the execution of contracts. There exist two main agents: the "customer" agent that needs resources from outside and bids for contracts and the "supplier" agent that provides services. The supplier agent offers its services to the customer agent. The customer agent in turn bids for the contract based on the cost, risk, and time constraints. A list of best bids is generated, the best bid is selected, and the execution manager (part of customer agent) is notified, which in turn oversees the contract to its completion.

6.3 Reconnaissance Agents

Reconnaissance agents (RAs) are programs that look ahead of user's browsing activity and act as a recommendation tool for the user to select the most appropriate links and ignore the needless paths [9]. They look ahead of the user's link for the "404 Not Found" error or for the relevance of the user's search. The RA can warn the user of the irrelevance of the link or draw attention to a more relevant link. Libermann and others [9] have suggested two RAs: "Letizia" for local reconnaissance and "Powerscout" for global reconnaissance.

6.4 IAs for Citation Finding on the World Wide Web

For the scientific community searching for relevant papers, research is an important task. On-line citations having links to a research paper are valuable sources of information. Searching for citations through search engines on the Web does not yield satisfactory results. Loke and others [12] have suggested a method to develop an intelligent agent for finding citations on the World Wide Web (WWW). They search the Web for citations based on the author's name and the title of the paper. They provide many search strategies, such as searching the Web for a hypertext markup language (HTML) version of the paper, searching the author's home page, searching the university home page, or searching the technical-report archive. They query search engines to obtain starting points for their searches. They also provide navigational rules to follow further links (searches) in a page.

6.5 Web-Browser Intelligence

The WWW consists of two tightly coupled systems: Web browsers and Web servers. The user requests information on the WWW from browsers in the form of universal resource locators (URLs), links, and forms, and servers respond with a page appropriate to the user's request. This tight coupling has undesirable consequences-namely that the displays that Web servers generate are impersonal presentations that many Web users find unattractive and difficult to use. Barrett and others [6] have introduced a model, Web-browser intelligence (WBI), where a programmable intermediary is introduced between the browser and the server. WBI interprets browser commands and dynamically reconstructs the document returned by the server for viewing. The advantage of WBI is it can modify the appearance of the Web pages returned by the servers for more elegance. It adds colors around links to indicate and warn the user about the speed and traffic at the server. WBI can also maintain a "user requests" history and provide methods to query this history, generate popular pages, and recommend shortcuts to links.

WBI has a set of four main modules: monitors, editors, generators, and rules. Monitors record the list of URL links requested, which is useful for maintaining a request history. Editors change the content and appearance of Web pages by adding links or extra buttons. Generators handle Web requests, such as performing searches on a person's history based on a keyword. Rules define a list of specific monitors, editors, and generators to be used for a given request, such as using HTML-modifying editors only.

7. Conclusion

Significant research is being conducted in the area of IAs. They show promise in their usefulness and do not have significant disadvantages or performance hindrance. IAs have found their presence in many applications on the Web. Many models and tools are being developed at a fast pace, and new ideas are being researched. Not long from now, the IA's presence is going to be felt in many areas and applications.

8. Reference

[1] M. Wooldridge and N. Jennings, "Intelligent Agents: Theory and Practice." Submitted to Knowledge Engineering Review.

[2] E. Kendall et al., "An Application Framework for Intelligent and Mobile Agents," ACM Computing Surveys, Vol. 32, No. 1 (March 2000).

[3] S. Wai et al., "An Intelligent Agent for Citation Finding on the World Wide Web," Technical Report 96/4 (Dept of Computer Science, University of Melbourne, Australia).

[4] N. Muler, "Improving Network Operations with Intelligent Agents," International Journal of Network Management, Vol. 7, 116-126.

[5] "Intelligent Agents for Electronic Commerce," http://www.cs.umn.edu/Research/airvl/magnet/index.html (Accessed October 29, 2001).

[6] R. Barrett et al., "Web Browser Intelligence: Opening Up the Web," IEEE proceedings of COMPCON '97.

[7] S. Shearin and H. Lieberman, "Intelligent Profiling by Example." Proceedings of the International Conference on Intelligent User Interfaces (2001).

[8] D. Lange and M. Oshima, "Seven Good Reasons for Mobile Agents," Communications of the ACM, Vol. 42, No. 3 (March 1999).

[9] H. Lieberman et al., "Exploring the Web with Reconnaissance Agents," Communications of the ACM, Vol. 44, No. 8 (August 2001).

[10] S. Franklin and A Graesser, "Is It an Agent, or Just a Program?: A Taxonomy for Autonomous Agents." Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (Springer-Verlag, 1996).

[11] Brustoloni and C. Jose, "Autonomous Agents: Characterization and Requirements," Carnegie Mellon Technical Report CMU-CS-91-204 (Pittsburgh: Carnegie Mellon University).

[12] S. Loke et al., "An Intelligent Agent for Citation Finding on the World Wide Web," Technical Report 94/4 (Department of Computer Science, The University of Melbore, Australia).

[13] O. Etzioni et al., "Intelligent Agents on the Internet: Fact, Fiction, and Forecast," IEEE Expert, Vol. 10, No. 3 (1995), 44-49.

[14] M. Grimley and B. Monroe, "Protecting the Integrity of Agents: An Exploration into Letting Agents Loose in an Unpredictable World," ACM Crossroads, Intelligent Agents (Summer 1999), 5.4.

Restoration in Layered Architectures with a WDM Mesh Optical Layer

Krishna Bala
Chief Technical Officer
Tellium Inc.

Eric Bouillet
Senior Network Architect
Tellium Inc.

Sid Chaudhuri
Director, Network Architecture
Tellium Inc.

Georgios Ellinas
Senior Network Architect
Tellium Inc.

Jean-François Labourdette
Manager, Algorithms and Performance Analysis Group
Tellium Inc.

Ramu Ramamurthy
Senior Network Architect
Tellium Inc.

Abstract

This work provides an overview and a comparison of various techniques used to restore lightpaths in a layered architecture. The comparison takes into consideration figures of merit such as speed of restoration and redundant capacity. The objective is to optimize the performance of the network (including fast end-to-end service survivability, minimizing the spare resources required) under all circumstances, using available resources and implemented mechanisms.

1. Introduction

It is widely accepted that optical cross-connects (OXCs) will be used to implement the next-generation mesh optical networks [1]. Optical network equipment vendors are currently implementing the next generation of optical switching systems capable of switching hundreds of lightpaths, each carrying millions of voice calls or thousands of video streams. Optical network architectures not only provide transmission capacities to higher transport levels, but also the intelligence required for fast lightpath provisioning and fast and efficient failure restoration [2,3,4]. The emergence of intelligent optical network elements is instrumental in making such optical architectures a reality today.

Optical network architectures can suffer failures and breakdowns, either due to accidental fiber cuts and operation mistakes or equipment malfunctions, such as switch, card, or component failures. Two competing approaches are being proposed for providing the appropriate recovery mechanisms in these circumstances. In the *peer-to-peer* approach [5,6,7], interweaved optical and higher-layer equipment act in symbiosis under the same control plane. In the *overlay* approach [5,6], optical and higher-layer domains are two separate entities with individual control planes, exchanging management services through a standard interface. The peer-to-peer approach relies on a unified bandwidth-management protocol to reassign bandwidth away from defective areas in the network and re-establish the interrupted data services. In the overlay approach, each layer independently relies on its own restoration mechanism in a manner that is independent and transparent to one another.

A single network layer or a combination of multiple layers can be used for failure restoration in a layered architecture. The aim is to provide service protection against a variety of failure conditions while restoring all failures quickly and with the minimum amount of protection capacity. Failure restoration involving multiple layers can enhance end-to-end survivability of the service by having each layer's protection scheme supplement each other. It is important to note, however, that multilayer protection may not be required or may be difficult to implement because of race conditions and complex escalation strategies and inter-layer protocols.

Section 2 presents the layered network architecture and describes each layer. Section 3 examines how restoration can be addressed in such anarchitecture, and Section 4 addresses restoration specifically in the optical layer. Section

5 discusses whether multilayer restoration is necessary and addresses possible escalation approaches in a network architecture where multiple layers can be used to restore a service. Concluding remarks follow in Section 6.

2. Overview of a Layered Network Architecture

In this section, we review the fundamental parts that constitute a network and its functionality. It goes without saying that many architectures exist or have been suggested, and it is not the intention of this paper to enumerate them exhaustively (see [5,8,9,10] for further information and useful references on this topic). However, we observe that all of the proposed architectures repose on a common denominator. It is this generic model that we present here. The model consists of three superimposed layers. Each layer provides well-defined services to its superjacent layer while concealing implementation details to it. As shown in *Figure 1*, from top to bottom, the layers are the service layer, the logical (electrical) layer, and the optical layer.

2.1. Service Layer

In the service layer, clients such as edge or service routers or MSPPs located in a provider point-of-presence (POP) represent users and the data communications among them. Using a graph representation, a node corresponds to a client that emits and receives data, and a link represents a service or a two-way data stream between clients. Link attributes in this layer correspond to minimum quality of service (QoS) requirements, which transpose into bit rates, jitter, and bit-propagation or round-trip delay constraints. Service-level agreements (SLAs), for instance, are negotiated and crafted in this layer.

2.2. Logical Layer

Also known as the electrical or digital layer, the logical layer aggregates services into large transmission "pipes" and assures their proper routing from POP to POP with prescribed QoS. Using a graph representation, a logical node corresponds, for example, to an Internet protocol (IP) core router, an asynchronous transfer mode (ATM) backbone switch, or a digital cross-connect, and a logical link connects the ports of two adjacent nodes. Capacity of a link and data processed by the nodes are expressed in units of bits per second (bps) in increments of digital signal (DS)–3 (45 Mbps) to optical carrier (OC)–192 (10 Gbps). The logical layer may consist of several interconnected sub-networks, either for scalability reasons, as it is easier to manage several smaller networks than a large network (hierarchical decomposition), or because the sub-networks belong to several independent carriers (multivendors, two-tier networking) or employ different technologies (e.g., IP versus ATM). In either case, boundaries and proper network interfaces within the logical layer delimit the sub-networks and their respective domains of operation.

The logical layer fulfils several roles: it maintains a consistent topological view of the layer; it manages the address space; it routes streams on request; and it polices the traffic to ensure a fair share of capacity among data streams and to guarantee each individual's QoS. The first part, also called topology discovery, is achieved by way of a neighbor discovery protocol (NDP) in conjunction with the open shortest path first (OSPF) protocol[1]. NDP operates in a distributed manner through in-band signaling to construct local port-to-port connectivity databases at each node. OSPF completes the topology discovery by assembling and globally disseminating pieces of information collected by NDP, plus additional information such as link states, to the logical plane. Logical nodes have only a few tens of ports, and with the exception of very small networks, a full connectivity featuring one link between every pair of node is not probable. Instead, services may have to be routed in the logical layer through one or more transit nodes to the desired destination using the constraint-based routing–label distribution protocol (CR–LDP)/resource reservation protocol (RSVP) explicit routing and bandwidth reservation protocols. The computa-

FIGURE 1

Layered Architecture

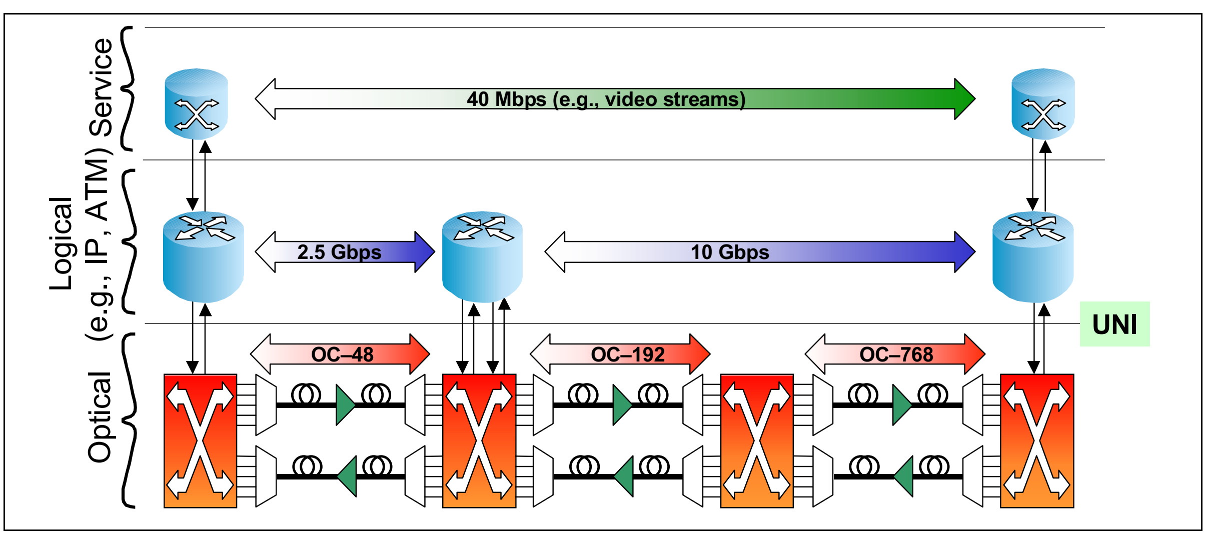

tion of a logical path must satisfy a set of constraints, such as round-trip delays and spare bandwidth, defined in the service layer in accordance to prescribed QoS. Note that the failure of a logical link or logical node is detected by NDP, and advertised by OSPF. That is, the layer has the primitives to detect a failure and resume interrupted services.

2.3. Optical Layer

The optical layer offers and manages the capacity required to transport traffic between clients in the logical layer. *Figure 2* depicts an example of a logical network (two IP routers) linked to an optical network (four optical switches). Optical switch ports are either add/drop ports interfacing the optical layer to the client's logical layer or network ports interconnecting optical switches. Using our graph representation, nodes are optical switches, and links are bundles of bidirectional optical channels between pairs of optical switches. An optical channel is a wavelength that connects the network ports of adjacent optical switches. A link in the logical layer is realized by way of optical channels in tandem forming a lightpath (circuit) between the end nodes of that link.

The optical layer faces the same challenges, and conceptually even borrows solutions from the logical layer. For instance, it relies on generalized MPLS (GMPLS) [11–13], also formerly known as MPLambdaS (an extension of MPLS) to encompass all types of architectures, including wavelength-oriented traffic engineering and management. It also relies on the NDP/link-management protocol (LMP) [6,7,14] and OSPF protocol [15] to create and publicize the network's topological views. Differences that set apart the optical layer from its logical counterpart are, among others, as follows: routing in the optical layer is exclusively circuit-oriented; circuit set-up and tear-down is done at a much slower time scale; and the bandwidth granularity of the logical layer is much lower than the granularity of the optical layer.

In the overlay approach, the layers work individually, with the client logical layer leasing resources from the optical layer. The user network interface (UNI) harmonizes communication of control messages between the two domains. The Optical Interworking Forum (OIF) is currently specifying UNI requirements and is working on their standardization [16]. In addition, an optical carrier will normally acquire network components from several vendors. A suite of protocols is being developed in the Internet Engineering Task Force (IETF) to allow for the seamless interaction between the various network components. As part of the suite, the LMP, for example, is used to maintain control channel connectivity, verify component link connectivity, and isolate link, fiber, or channel failures within the network [17].

3. Failure Restoration in a Layered Architecture

Restoration is invoked upon failure of one or more elements along paths carrying end-to-end services. Both logical and optical layers may implement an autonomous recovery scheme, and both may react to the same defect. Failures are either one of two types: logical (such as a malfunctioning IP router) or optical (for instance, a fiber cut). There are thus four possible scenarios, depending on the origin of the failure and the layer that provides the restoration [18]:

1. Failure and restoration in the optical layer as shown in *Figure 3(b)* based on the original routing shown *in Figure 3(a)*. *Figure 3(a)* illustrates the connectivity in the optical layer and the resulting connectivity in the logical layer during normal mode of operation. *Figure 3(b)* is an example of optical failure restored in the optical layer. The affected lightpath is restored away from the failure using optical capacity that was reserved for this purpose. The operation is transparent from the logical layer, which remains unchanged.
2. Failure and restoration in the logical layer as shown in *Figure 3(c)*. This figure illustrates a logical failure (ATM switch or IP router failure) restored in the logical layer. After failure, the service is rerouted using the remaining capacity of the logical layer. The operation is transparent to the optical layer.

FIGURE 2

Optical Layer

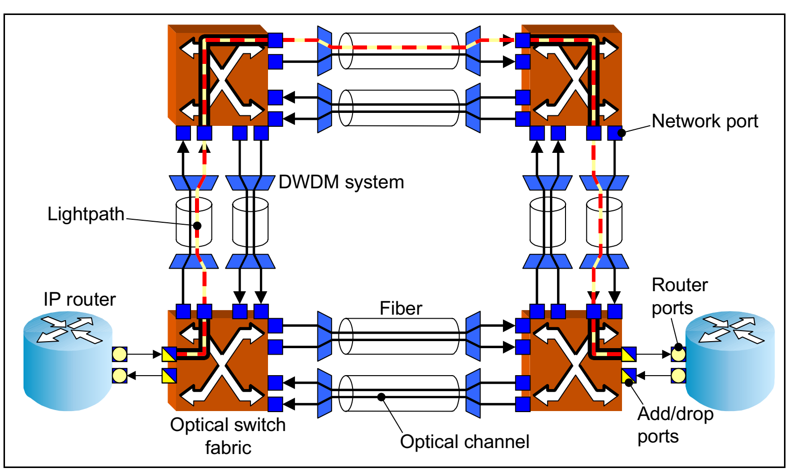

3. Optical failure repaired in the logical layer as shown in *Figure 3(d)*. This figure illustrates an optical failure restored in the logical layer. If the optical layer fails to restore the optical failure after a certain time lapse, the logical layer can restore the service on a different logical path, using for instance implicit LSP protection in MPLS.
4. Logical failure repaired in the optical layer as shown in *Figure 3(e)*. Unlike any of the previous protection schemes, restoring a logical failure with leverage from the optical layer involves reconfiguration with creation of new connections in the logical layer and the optical layer. This type of restoration may be necessary if, after logical failure, the remaining capacity in the logical layer is insufficient to reroute all of the affected services. Additional logical capacity can be created with the provisioning of new lightpaths.

A fifth and most realistic situation consists of optical failures repaired simultaneously and independently in both layers. Since this is a combination of the aforementioned scenarios, it is not considered here.

Restoration in an IP–centric logical layer is accomplished by MPLS [19]. MPLS enables a hierarchy of label-switched paths (LSPs) to be defined by pre-pending a stack of labels or tags to packet headers. Upon an optical or electrical failure occurrence, packets along a given disrupted LSP can be routed to a predefined restoration LSP by modifying the label maps of the routers at the end-points of the original LSP [20]. In a similar manner, restoration of optical failures in the optical layer is also achieved by way of redundancy. Studies indicate that restoration in the optical layer requires substantially more spare capacity, depending on the diligence and the quality of the protection, yet overall the solution is more economical due to a lower cost per units of capacity [20].

MPLS offers undeniable potentials for fast restoration. The principal advantage of MPLS is its ability to recover indiscriminately from failures in the logical layer or the optical layer, as suggested in *Figure 3(c)* and *Figure 3(d)*. However, a single failure may affect thousands of LSPs and trigger an avalanche of alarms and corrective actions. The resulting

FIGURE 3

(a) Routing before a Failure Occurs; (b) Failure and Restoration in the Optical Layer; (c) Failure and Restoration in the Logical Layer; (d) Optical Failure Restoration in the Logical Layer; (e) Logical Failure Restoration in the Optical Layer

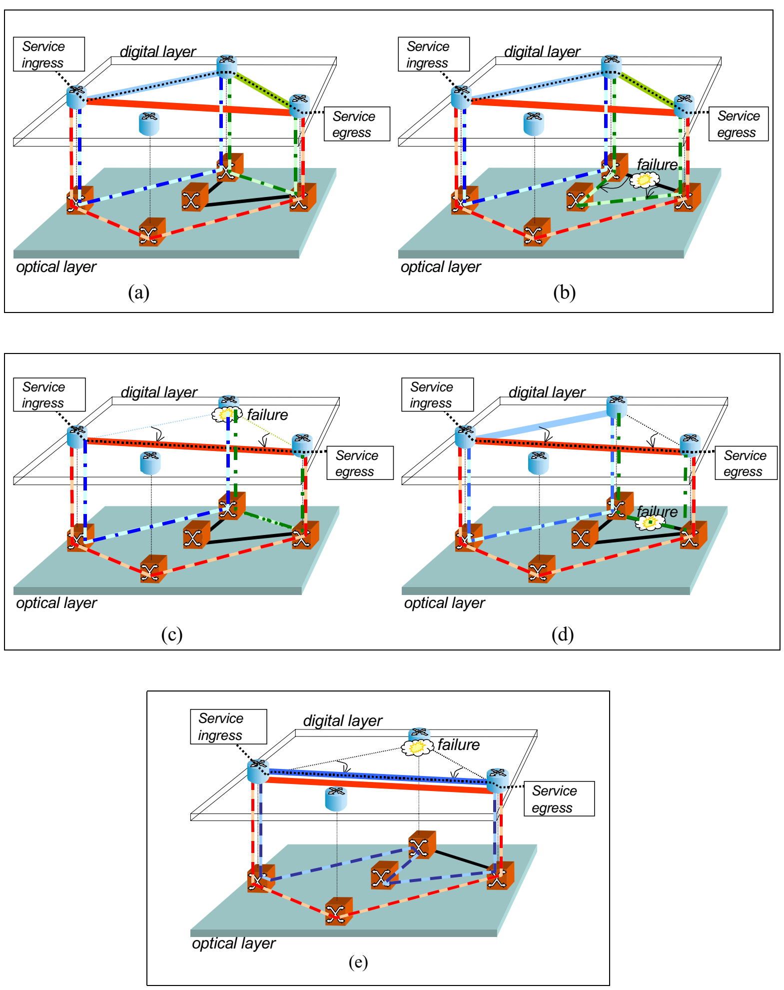

amount of signaling can be orders of magnitude higher than in the optical layer, which is able to switch hundreds of LSPs multiplexed into a single wavelength at once. Also in MPLS restoration, primary and back-up LSPs must not succumb together to a malfunction in the logical or in the optical layer. Tto satisfy the second condition, the logical layer must explicitly inquire about the risk relationship between the lightpaths that compose its logical connectivity and compute the LSPs, primaries, and respective back-ups accordingly. Proposed specifications for the UNI interface allow the logical layer to request lightpaths that are disjoint from selected subsets of pre-established lightpaths. However, this approach yields lower availability than other approaches that allow the optical layer to decide on the restoration mechanisms. Such requests may thus sometime be impossible to realize, even if the capacity is available in the optical layer. Another strategy is to rely on NDP, OSPF, and IP self-routing properties to advertise and correct failures in the logical configuration, but then the restoration time is not as attractive in terms of restoration speed as it would be with predefined restoration LSPs.

The fourth scenario implies a minimum of synergy between the restoration architectures deployed in each layer; the optical layer does not know a priori the logical connectivity of the client, and hence cannot take the initiative to restore a logical failure. Both layers, however, could coordinate their effort to resume interrupted services, with the optical layer getting directives from the logical layer. In particular, the logical layer could provision spare capacity in the optical domain and reclaim some of it upon failure of one of the routers to create new logical connections and balance the load on the surviving routers. The feasibility of this scheme is subject to UNI specifications.

To summarize, although the aforementioned scenarios one and three address the same problem of recovering from failure in the optical layer, the first, which recovers the failure in the layer where it occurs, is preferable in terms of cost and speed [21,22]. The same is also true with the second over the fourth scenario. In addition, because the preferred mechanisms are confined within their own layers, that helps to simplify the restoration approach and to avoid architectural complexities and the interdependence of mixed-layer approaches.

4. Restoration in the Optical Layer

In end-to-end OXC–based path protection, the ingress and egress nodes of the failed optical connection attempt to restore the signal on a predefined back-up path, which is link-disjoint from the primary path. Path diversity guarantees that primary and back-up lightpaths will not simultaneously succumb to a single failure. There are two sub-types of path protection: 1+1 dedicated protection and mesh restoration.

4.1. Dedicated Protection

Dedicated 1+1 protection is illustrated in *Figure 4*. The network consists of four logical nodes (A to D) and two demands (AB and CD) accommodated across an eight node optical network (S to Z.) The provisioning algorithm of this architecture computes and establishes simultaneously the primaries and their link-disjoint protection paths. During normal operation mode, both paths carry the optical signal, and the egress node selects one of the two copies. In the example of *Figure 4*, all of the optical channels on primary and secondary paths are active. In particular, the configuration reserves two optical channels between nodes S and T for protection. This is the fastest restoration scheme, since, for every lightpath, one device at the termination of the lightpath is responsible for all of the necessary failure detection and restoration functions. But it is also the most exigent in terms of resource consumption.

4.2. Mesh Restored Lightpaths

As in dedicated protection, in mesh restoration back-up paths are predefined, but the cross-connections along these paths are not created until a failure occurs. During normal operation modes, the spare optical channels reserved for protection are not used. Because the capacity is only "soft reserved," the same optical channel can be

FIGURE 4

Dedicated 1+1 Protection

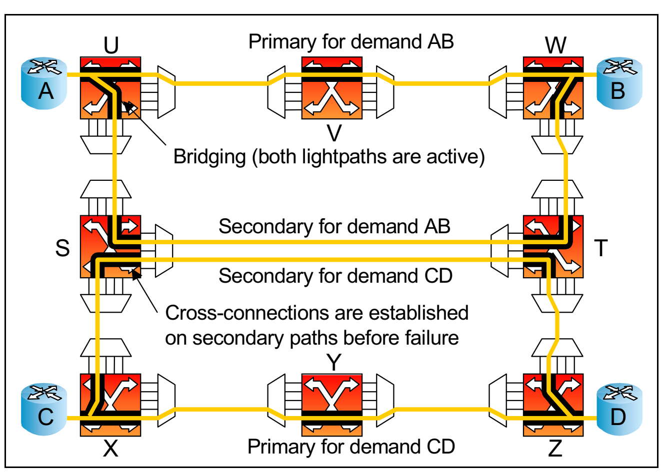

shared to protect multiple lightpaths. There is a condition, though, that two back-up lightpaths may share a reserved channel only if their respective primaries are link-disjoint, so that a failure does not interrupt both primary paths. If that happened, there would be contention for the reserved channel, and only one of the two lightpaths would be successfully restored. *Figure 5(a)* (for normal mode) and *Figure 5(b)* (for restoration mode) picture an example of mesh restoration. The dashed lines represent reserved channels. In this case, the protection paths for demands AB and CD share a single optical channel in link S-T, one less than in dedicated protection. However, the restoration involves a bit more processing to signal and to establish the cross-connections along the restoration path. There is thus an evident trade-off between capacity utilization and recovery time.

Simulation experiments were run on a 100-node (N100), 137-edge network that has a degree distribution of (50, 28, 20, 2) nodes with respective degrees (2, 3, 4, 5). It was assumed that this architecture has infinite link capacity. Three network robustness scenarios were considered: no protection, dedicated protection with provisioning to recover from single link failures, and mesh restoration to recover from link failure. In N100, 3,278 node pairs out of 4,950 possible node pairs are connected by one bidirectional lightpath. Requests for lightpaths arrive one at the time (on-line routing) in a finite sequence and in an order that is arbitrary but common to each scenario to ensure a fair comparison. Figures of merit are capacity requirements separated into their primary and restoration parts and expressed in units of bidirectional OC–48 channels. Results are presented in *Figure 6*. The quantities shown in the chart are averages over a series of 10 experiments using various demand arrival orders. These results clearly demonstrate the advantages of shared mesh restoration over dedicated protection in terms of redundant capacity required to protect against all single link failures [23].

5. Multilayer Protection

Notwithstanding the architectural complexities and interdependence of mixed-layer restoration approaches, this section investigates how more than one layer can offer protection in a multilayer network. The goal of such an approach would be to use the protection capabilities of each layer to provide additional survivability capabilities in the network.

A number of different multilayer protection strategies can be implemented in a layered architecture. As pointed out in the previous section, restoring the network at the layer where the failure occurred is a fast approach better suited for the recovery of failures in that layer. In addition, protection and restoration in different layers may be mixed together for the best overall result. For instance, fast optical protection architecture for fiber and OXC failures can be supplemented by service-based restoration at the logical layer. In this case, the optical layer can offer bulk recovery of the services while the logical layer can offer finer restoration granularity.

5.1. Escalation Strategies

If a multilayer protection approach is adopted, an escalation strategy has to be provided to coordinate the protection processes of the different layers. The absence of an escalation strategy can create race conditions between the protection mechanisms with unpredictable (and potentially catastrophic) results. This section identifies the issues associated with the escalation strategies and presents a comparison of the different escalation options.

The escalation strategies can include either a parallel or a sequential activation of restoration mechanisms. In the parallel approach, restoration mechanisms from different layers are trying to restore the same failure simultaneously, which will result in a very fast restoration time. However, the different restoration mechanisms must be coordinated so as not to obstruct each other or compete for the same restora-

FIGURE 5

Mesh Restoration (a) Network Connections before a Failure Occurs; (b) Network Connections after a Failure Occurs

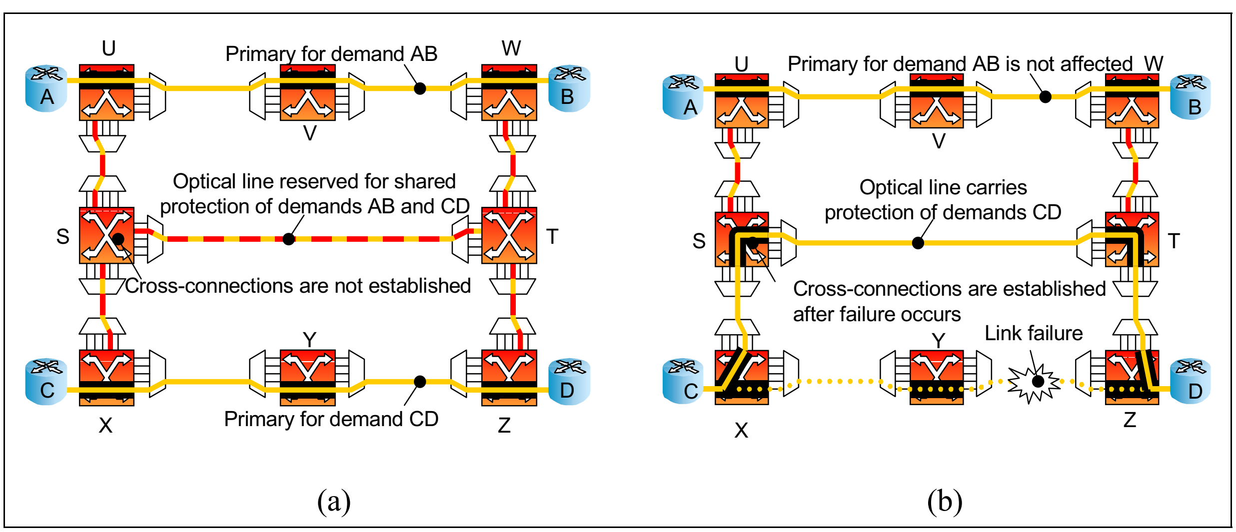

tion resources. In the sequential case, restoration mechanisms from different layers attempt to restore the failure one layer at a time. One sequential approach could be to wait until one layer has failed to restore the services or a fixed time interval has passed before the restoration process is taken over by another layer [24,25].

Typically, the sequential approach will be slower in terms of overall restoration times, but it is more easily implemented. Such a strategy only needs to predetermine the order in which the layers will attempt restoration and when the transition from one layer to another takes place. The type of failure will usually determine the layer where the sequential restoration process will start. For example, if a fiber is cut, restoration can start at the optical layer, where it can achieve fast restoration of bulk traffic. On the other hand, if an IP router fails, restoration can start at the logical layer. This will allow for finer granularity of the restoration process, but it will be much slower.

A hold-off timer function can be used in the sequential approach to mark the transition from the protection mechanism of one layer to the protection mechanism of another layer. Although this is a simple approach, it can result in a cascade effect and potentially introduce considerable delay to the restoration process. Another sequential escalation strategy can use a diagnostics method to ascertain that a protection mechanism will not be successful before the hold-off timer expires and subsequently hand over the restoration control to the next layer. Even though this scheme will result in much faster restoration times, this is a much more complex approach to implement. Other escalation strategies include signaling approaches (e.g., through the UNI or through the management system) to coordinate the passing of the protection responsibilities to another layer.

Take for example a multilayer network consisting of IP, Gigabit Ethernet (GbE), and wavelength division multiplexing (WDM) layers. Restoration in such a network can take place at the WDM and/or IP layers, as GbE does not support restoration capabilities. If the WDM layer only is used for restoration, fast restoration can be utilized to quickly restore services during a network-failure condition at the optical layer. If the IP layer only is used for restoration, router-based rerouting will result in longer restoration times, but it will require less redundant capacity because of the finer restoration granularity in this layer. In general, using the IP layer only for restoration may not meet the service's QoS requirements. However, for some data services, such as e-mail and file transfer protocol (FTP), restoration at the IP layer only may be sufficient.

If a two-layer strategy is employed, a parallel escalation strategy may suffice. Because of the different time scales between the restoration mechanisms in the WDM and IP layers, the WDM layer will restore the traffic before the IP layer even detects the error. A sequential approach on the other hand would require coordination between the optical layer and the routing protocols used in the IP router-based rerouting. The best approach for such a network would be to use the optical layer for the restoration of time-sensitive services and use IP router-based rerouting for the restoration of less critical and time-insensitive services.

6. Conclusion

This paper describes multilayer networks and how a single layer utilizes its protection capabilities to restore services after a failure occurs. Specifically, restoration in the optical layer is described, and dedicated and mesh lightpath-based restorations are presented. A number of escalation strategies are also presented for the case in which a combination of different restoration methods from different layers is to be used. From the discussion, it is evident that the simplest solution is to allow the layer where the failure occurred to restore for that failure. This will help in avoiding functional

FIGURE 6

Comparison of Capacity Usage for Different Protection Architectures (100 Nodes)

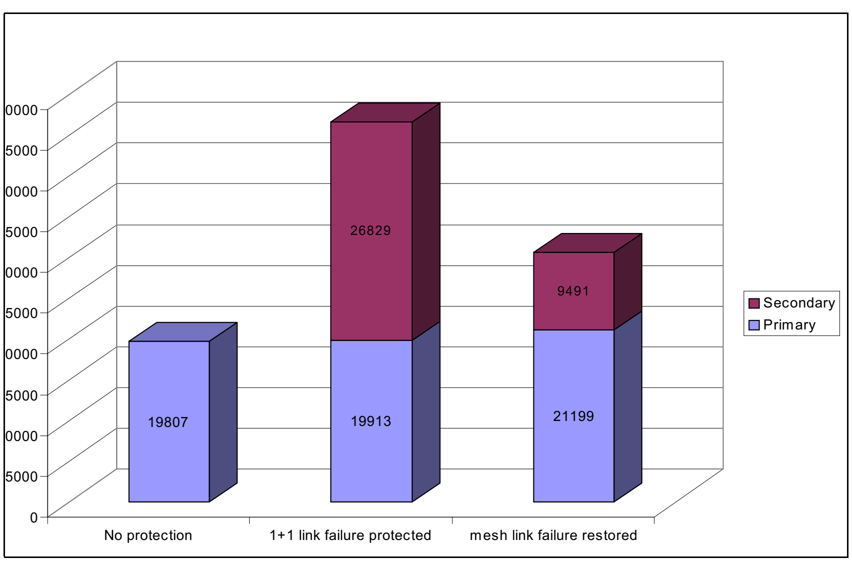

duplication, race conditions, as well as complex interaction and inter-layer coordination problems. Even though allowing for a multilayer restoration mechanism can enhance the overall survivability of the network, complex escalation techniques will violate the layering principle by violating inter-layer independence. Because the optical layer is the lowest layer in the transport hierarchy, using the restoration mechanism in that layer is ideal for fast recovery of critical services during failure conditions such as fiber cuts and optical switch failures. In contrast, IP router-based rerouting is a best-effort approach more suitable to the restoration of less critical and time-insensitive data.

References

[1] T.E. Stern, K. Bala, *Multiwavelength Optical Networks: A Layered Approach* (Prentice Hall, May 1999).

[2] D. Benjamin, R. Trudel, S. Shew, and E. Kus, "Optical Services over an Intelligent Optical Network," *IEEE Communications Magazine*, September 2001, 73–78.

[3] E. Varma et al., "Architecting the Services Optical Network," *IEEE Communications Magazine*, September 2001, 80–87.

[4] A. McGuire, S. Mirza, and D. Freeland, "Application of Control Plane Technology to Dynamic Configuration Management," *IEEE Communications Magazine*, September 2001, 94–99.

[5] "IP over Optical Networks: A Framework," draft-ietf-ipo-framework-00.txt, IETF draft, January 2001.

[6] B. Rajagopalan et al., "IP over Optical Networks: Architectural Aspects," *IEEE Communications Magazine*, September 2000.

[7] G. Bernstein, J. Yates, and D. Saha, "Control and Management of Optical Transport Networks," *IEEE Communications Magazine*, October 2000.

[8] G. Hjálmt_sson, J. Yates, S. Chaudhuri, and A. Greenberg, "Smart Routers—Simple Optics: An Architecture for the Optical Internet," IEEE/OSA Journal of Lightwave Technology, December 2000.

[9] G. Bernstein, J. Yates, and D. Saha, "IP-Centric Control and Management of Optical Transport Networks," *IEEE Communications Magazine*, October 2000, 161–167.

[10] K. Sato and S. Okamoto, "Photonic Transport Technologies to Create Robust Backbone Networks," *IEEE Communications Magazine*, August 1999, 78–87.

[11] "Generalized Multi-Protocol Label Switching Architecture," draft-many-gmpls-architecture-00.txt, IETF draft, August 2001.

[12] "Generalized MPLS—Signaling Functional Description," draft-ietf-mpls-generalized-signaling-05.txt, IETF draft, January 2001.

[13] "Generalized MPLS Signaling—RSVP-TE Extensions," draft-ietf-mpls-generalized-rsvp-te-04.txt, IETF draft, January 2001.

[14] "Link Management Protocol," draft-ietf-mpls-lmp-02.txt, IETF draft, September 2001.

[15] "OSPF version 2," IETF RFC 2328, April 1998.

[16] "User Network Interface (UNI) 1.0 Signaling Specification," OIF contribution OIF2000.125.6, September 2001.

[17] "Link Management Protocol (LMP) for WDM Transmission Systems," IETF draft, June 2001.

[18] P. Demeester at al., "Resilience in Multilayer Networks," *IEEE Communications Magazine*, August 1999, 70–76.

[19] S. Makam et al., "Framework for MPLS Based Recovery," draft-ietf-mpls-recovery-frmwrk-03.txt, IETF draft, July 2001.

[20] R. D. Doverspike et al., "Transport Network Architectures in an IP World," Infocom 2000 Proceedings (2000), 305–314.

[21] J. Manchester, P. Bonenfant, and C. Newton, "The Evolution of Transport Network Survivability," *IEEE Communication Magazine*, August 1999, 44–51.

[22] R. Doverspike, S. Phillips, and J. Westbrook, "Transport Network Architectures in an IP World," Infocom 2000 Proceedings (2000).

[23] E. Bouillet, et al., "Routing and Restoration Architectures in Mesh Optical Networks," to appear as an invited paper in the *Optical Networks Magazine*.

[24] "Planning for Service Survivability in Broadband Multilayer Networks," Bellcore SR-4317, August 1997.

[25] R. Batchellor, "Coordinating Protection in Multiple Layers", OIF Contribution, oif99.038.0, April 1999.

Notes

1. In this paper, the assumption is of a single network running IP–centric multiprotocol label switching (MPLS) protocols.

Metronomics

Scott Berry

Director

Metromedia Fiber Network

This paper discusses metronomics, or the economics of bandwidth delivery in the metro, particularly via fiber. Consideration will be given to the market, architectures, fiber, and the new networks and infrastructure in the metro going forward.

Bandwidth Distribution

Bandwidth is not merely growing—there is a fundamental shift occurring in the distribution of bandwidth demand. In recent years a substantial amount of bandwidth demand has been driven by corporate data needs. The Internet is still relatively small, but it is growing rapidly; voice is fairly large, but it is shrinking. The projections for 2005 indicate significant Internet growth, although as is currently the case a large portion of this growth will be corporate driven.

From another perspective, newer applications require more and more bandwidth—and by extension higher-speed connections—in much the same way that software applications have driven the development of higher-speed processors. The result is an exponential growth in bandwidth demand per application, something likely to continue as video and other high-speed applications take hold.

The primary enterprise application driving this exponential growth right now is data storage. As companies are becoming more Internet- and data-centric, a larger and larger part of what they produce is going to be data—and this must be stored. But as the amount of inexpensive bandwidth available to the corporation increases, a number of other applications become possible that weren't before, and some of these applications will be revenue generating. Thus the drop in bandwidth pricing not only improves an enterprises' bottom line; it can grow the top line as well. While storage is the current "killer app" for metro bandwidth, video (supporting such applications at telemedicine and distance learning) may well be next in line.

The most popular approach to network architecture has been to build "big pipes" in the long-haul backbone. However, in reality all data starts and ends with the end user, and so any network that does not connect directly with the data sources and sinks will not capture the bulk of the bandwidth demand. You can't have a nervous system built on a spinal cord alone. Because of the large amounts of bandwidth being required by enterprises, this connectivity will, by necessity, be fiber based. Customers are connecting with fiber to regional offices, to hosting centers, and from there to the Internet. Just as voice drove infrastructure in the past, data will drive it in the future. Today's Internet hosting centers will become the central offices of tomorrow.

Contrary to the perception in some quarters of a fiber and bandwidth glut, such a situation cannot arise until each part of the network has a glut, and that's nowhere near the case today.

The key advantage of fiber, of course, is that its capacity is virtually unlimited; it is determined only by the equipment placed on its ends—and technological advance in optical end-gear is happening at exponential rates. Fiber thus provides future-proof infrastructure. Although better fiber is being manufactured all the time, what is going in the ground today can hold many orders of magnitude more bandwidth than is currently demanded by the typical enterprise. In addition, extra conduits can be placed when fiber is installed to permit more fiber to be pulled as needed in the future.

Although bandwidth is generally priced by the bit, so that greater usage leads to higher costs for users, fiber is priced at a flat rate. This provides a very attractive financial model for enterprises and other high-bandwidth users because it not only cuts costs but also facilitates new, revenue-generating applications. Savings actually grow as bandwidth usage increases. Indeed, we have seen that as the price for bandwidth drops, more is consumed, indicating true elasticity of demand.

For example, in a typical metro area, procuring digital signal (DS)–3 service from a local provider might cost roughly $12,000 a month for some 200 megabits of capacity. Alternatively, several buildings could be connected with private fiber, and if synchronous optical network (SONET) gear were installed, and the cost amortized over three to five years, over 10 times the capacity might be realized at the same monthly cost as for traditional telco service. Moreover, adding the next DS–3 would be as simple as installing relatively inexpensive line cards and thus paid for only once, while additional capacity from a carrier would be paid for each and every month.

Several different options exist for customers to obtain bandwidth. As in the past, they can obtain carrier service. They can also lease dark fiber and then purchase optical equipment—dense wavelength division multiplexing (DWDM) equipment, a router with optical interfaces, or SONET gear, for example—and manage it independently. Customers can run their own networks much the way that companies run their internal networks. The last option is for customers to lease dark fiber and then go to a third party who offers a value-added service to equip and operate the fiber network for them.

Typically, the current architecture in the metro involves a service provider aggregating multiple customers' traffic out of a building into one or two central offices using public switches and public facilities. Private, fiber-based solutions, however, are different because they are unshared. Nobody else's traffic is on that fiber; the end equipment is utilized 100 percent by the customer and nobody else. Nor does the traffic on that fiber need to transit central offices or other shared aggregation points. It is, therefore, a truly private network. By using a third-party outsourcing method for managing that network, customers avoid a capital outlay for the equipment. They would have a monthly lease for fiber and a monthly charge for equipment, maintenance, and monitoring, all bundled as a single service. While this resembles the pricing model for telecom service, the cost can be orders of magnitude lower for the same capacity.

Take the case of a customer with six buildings connected in a metro area with two optical carrier (OC)–48 rings from a carrier (a total of about 5 gigabits). This service could easily cost $200,000 a month. With a fiber solution using DWDM equipment, for example, that customer would be able to do 24 gigabit-Ethernet channels (many times the capacity) between those same six buildings for less money.

Or, take a customer with a rapidly growing need for infrastructure. They have an interoffice local-area network (LAN), do a substantial amount of transaction processing with the financial services industry, and have six sites. The connectivity required is meshed gigabit Ethernet. With a fiber-based solution, they would be able to put one ring together with DWDM gear and provision six wavelengths to mesh those six gigabit Ethernet channels, fully protected with 50-millisecond switching. Plus they would have the flexibility to change the gigabit Ethernet to fiber channel or some other protocol at any time. This would be some 50 percent less expensive than what it paid for an equivalent telecom service, and the marginal cost to add another managed, wavelength-carrying gigabit-Ethernet channel would be about the price of a DS–3 on a monthly basis. The end result overall would be a twenty-fold improvement in its cost structure—plus the advantage of a truly private network.

These sample business cases also help demonstrate the price elasticity of bandwidth. In some instances, customers installing DWDM networks are actually ordering additional wavelengths before their networks are even up and running. Realizing how inexpensive it is, they are identifying a host of applications that had been thought infeasible because bandwidth was too expensive. In effect, the cheaper bandwidth becomes, the more of it they use.

Trends and the Future

Customers want flexibility in terms of where the traffic goes, as well as in terms of protocol and applications supported. Security and reliability are also critical. Enterprises do not want their data going through public switches if they can avoid it. Finally, they want the ability to use multiple (and non-standard) protocols—they are tired of being limited by the voice-driven SONET architectures of the past. Because of this, a private optical network is an appealing idea for many end-users.

Outsourcing is also a strong trend. Many customers may find the economics of fiber attractive, but they do not know how to run optical networks. And hiring engineers with the right talent is difficult, as is retaining them. In addition, once customers have fiber leaving the building, even if they are managing the optical gear themselves, they still have an outside vendor to deal with (namely the fiber provider) if there is a problem. In other words, you can't insource outside plant.

Finally, avoiding capital expense is an important benefit to such outsourcing. Rather than facing a significant investment up front to obtain less expensive bandwidth, customers can simply continue paying the same type of monthly bill they are used to, yet receive 10 to 20 times the bandwidth as before.

Demystifying eLearning Portals: The Convergence of Enterprise Intelligence and Learning

Bray J. Brockbank
Senior Vice President, Business, Technology, and Marketing
IBC, Inc.

Challenging Times

In recent years, the focus of organizational learning and training has shifted from traditional, new employee orientation and personal development seminars to continuous learning and development programs designed and implemented to provide employees with the knowledge resources necessary to successfully fulfill their roles and responsibilities. The goal of the new knowledge economy is to channel knowledge into ideas and use those ideas to create business solutions and competitive advantage. These ideas can then be channeled and used by the organization to thrive in an environment of intense competition, relentless change, and highly educated customers.

Organizations are burdened with channeling knowledge and ideas in the form of training and learning to employees more rapidly, more effectively, and in an even more efficient manner than ever. In an effort to meet these objectives, organizations are seeking to consolidate disparate knowledge and learning technologies into a centralized point of access and management, while leveraging existing investments (legacy systems, resources, etc.) in enterprise infrastructure in order to provide employees with simple, seamless access to knowledge, training, and learning.

Solid and intense competition is forcing organizations to implement collaborative solutions that integrate internal systems and leverage existing (legacy) technology to harness the knowledge and intellectual capital that exists in the public domain, while utilizing knowledge and intellectual capital that resides across the organization and among suppliers, partners, and customers.

Current Realities

The new economy is fueled by knowledge. Organizations possess incredible intellectual capital. The challenge has been, and continues to be, providing access to that capital and assembling it for the development of best practices and collective learning. Once defined marketplaces are melding into one global marketplace. Human capital is now an asset to be scrupulously managed. Learning is a strategic advantage and weapon. Workforce supply is in flux—everyone's a free agent. The world is in a state of rapid-growth and hyper-efficiency—and its all beginning to blur at Internet speed.

The critical and differentiating force of countries, organizations, and individuals lies in their knowledge and intelligence and how they use it in the new knowledge economy. eLearning is critical to the success of individuals, organizations, communities, and economies in the commencement of the knowledge economy. eLearning portals offer customization, 24/7 accessibility, convenience and flexibility, cost-effectiveness, just-in-time user-centric learning, and centralized management of knowledge.

In a recent IDC and eWorld survey (www.idc.com, IDC #24788, June 2001), training and education was rated in the top three applications integrated with U.S. organization Web sites. Nearly one-third of U.S. organizations have their training and education integrated with their Web sites. Training and education ranks third in priority (behind customer service and support and customer-relationship management [CRM]). More than 60 percent of large organizations are leading the way in integrating education and training with their Web sites—roughly the same as those who planned to integrate knowledge management and materials management into their Web sites.

Learning and training processes are becoming increasingly integrated into strategic organizational processes. Industries rapidly moving to integrate their training and education include health-care, government, education, banking, transportation, media, telecommunications, and utilities.

eLearning Industry

The eLearning industry is comprised of three vendor segments: technology, content, and services.

- *Technology:* This segment includes learning management systems (LMSs), learning content management systems (LCMSs), authoring tools, training delivery systems, enterprise resource planning (ERP), application service providers (ASPs), live eLearning tools, streaming video, EPSS, and testing and assessment tools.
- *Content:* This segment includes third-party content providers, books and magazine publishers, enterprises, subject-matter experts, government agencies, colleges, universities, schools, training organizations, eLearning portals, information technology (IT) firms, and system integrators.
- *Services:* This segment includes enterprise information portals (EIP), corporate universities, learning service providers (LSPs), content aggregators, learning consultants, consulting, professional services, certification service providers, collaboration services, and on-line mentoring services.

Even though many vendors are "pure-players," some vendors cross over or offer "hybrid" solutions. Several vendors market themselves as eLearning portals, end-to-end solutions, blended eLearning solutions, best-of-breed technology, global learning management solutions, integrated learning and management systems, and eLearning infrastructure technology.

Components of eLearning

eLearning components include LMS or LCMS, content, collaboration, testing and assessment, skills and competency, e-commerce, and Internet video-based learning (see *Figure 1*). A complete eLearning portal represents the total integration of multimedia, instructor-led, and real-time training—in a human, collaborative environment.

When implemented correctly, eLearning portals can help organizations to develop and maintain a competitive advantage in the following areas:

- *Recruitment and Selection:* The attraction, evaluation, and hiring of new employees
- *Retention:* The retention of intellectual capital (human capital)
- *Learning and Career Development:* Classroom training, on-line learning, and other forms of learning activities
- *Rewards, Recognition, and Response:* The recognition of individuals according to their ability to meet/exceed performance expectations as defined by the organization (according to appropriate business goals)
- *Succession Planning:* The identification and development of peak performers with the appropriate competencies and skills needed to advance within the organization

Types of eLearning Portals

eLearning portals can be categorized into five key categories: learning management, content aggregation, community collaboration, content creation, and internal.

- *LMS or LCMS Portal:* Tracks courses or learning content that the user is taking and his or her progress
- *Content Aggregation Portal:* Users can access knowledge on many topics from various sources
- *Community Collaboration Portal:* Users can interact with and learn from colleagues and peers
- *Content Creation Portal:* Where the user can create content directly on the site
- *Internal Portal:* Users can access all of the services of the organization's training and development department

Learning Management System

An LMS ties all other eLearning components together (see *Figure 2*). The LMS is the infrastructure or framework used to track, support, manage, and measure eLearning activities. An LMS helps to manage and measure the entire learning process, whether the organizations' needs are managing computer-based training (CBT), Web-based training (WBT), document-based training (DBT), instructor-led training (ILT), or blended training methods (BTM).

Learning Content Management System

An LCMS is a system used to create, accumulate, assemble, and distribute personalized eLearning (see *Figure 3*).

FIGURE 1

FIGURE 2

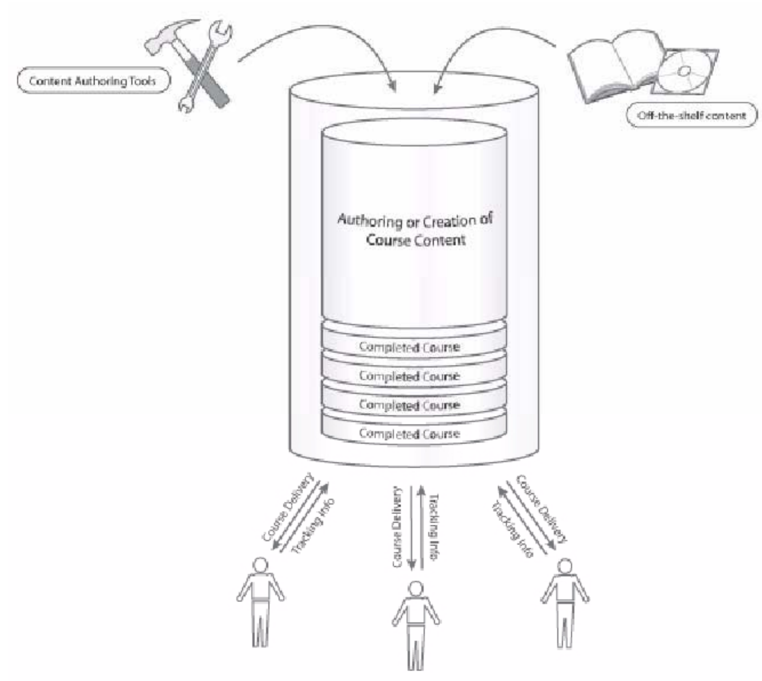

The LCMS was developed to more effectively deliver "small-chunks" of learning tailored to the users learning style and needs.

Content is delivered to the user in the form of learning objects. A learning object is an independent piece of education containing content and assessment based on specific

FIGURE 3

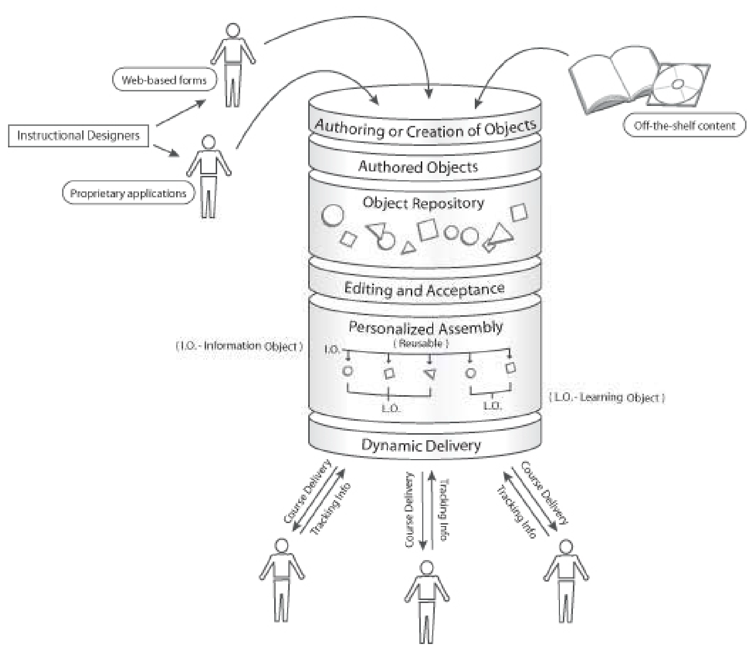

learning objectives. A learning object has descriptive metadata enclosed around it describing what is contained within the object. The object is assembled in an effort to help users to achieve specific learning goals and objectives.

Course Content

Organizations may assemble course content through various processes: buying off-the-shelf content, developing courses internally (in-house, home-grown), and outsourcing for custom courseware development. A complete eLearning portal should offer interoperability with all learning content.

- *Off-the-Shelf (or Third-Party):* Content developed to meet general organizational needs. Off-the-shelf content vendors include NETg, SmartForce, SkillSoft, etc.
- *Authoring Tools:* Authoring tool products include Authorware, Dreamweaver, etc.
- *Custom Developed:* Content developed to meet the unique needs of an organization and capture knowledge proprietary to a specific organization

Future-Built from the Ground Up

eLearning portals that offer scalability, flexibility, interoperability, and extendibility offer a "future-built" solution approach to eLearning.

- *Scalability:* How well the solution will work when the size of that need increases or decreases over time
- *Flexibility:* Allows the processing to be customized to meet business processes without writing additional code
- *Interoperability:* The ease with which the solution works and integrates with third-party content, collaboration, and software packages
- *Extendibility:* The capability of extending or adding new functionality to current or existing product features

Defining the eLearning Portal

There are basically two types of eLearning portals: external and internal. The first focuses on providing access to external learning services. The second focuses on providing access to all learning within the organization. Of course the internal portal may also include the use and management of external learning and providers.

Essentially, an eLearning portal is a virtual environment set up by an organization to give users access to knowledge. These portals have also been called eLearning centers, on-line education centers, internal portals, corporate universities, and virtual universities.

A portal is merely a vessel, framework, or infrastructure for training, learning, and the assessment of knowledge and competency. With the advent of eLearning portals, organizations now have tools to help knowledge workers aggregate, access, and navigate through full or bite-sized "learning chunks" or "learning objects" from internal databases, repositories, courses, and Web sites. The complete eLearning portal supports the learning cycle with various components of eLearning.

Market Categories

eLearning portals service three primary market categories: corporate, academic, and consumer. An eLearning portal is an access point to various applications, services, courses, and learning directed at one or more of three user groups: employees, partners/suppliers, and customers. The eLearning portal is a personalized, single point of access for the internal and external user where the organization's Web channels come together—the Internet, intranet, extranet, and marketplace exchanges—to exploit the cumulative information, knowledge, and data that will enable greater business efficiencies.

Intelligent Portals

The eLearning portal is an intelligent portal. The portal advises users on what skills and experience they need to advance to other levels in the organization, provide competency maps and assessment, and discussion forums related to essential learning themes—on-line learning communities. It recognizes what the user knows, certifications earned, experiences, and his or her ideal learning style. As eLearning portal technology has evolved, navigation has become more sophisticated, content more relevant, and interfaces more user-friendly and intuitive. Until most recently, eLearning was offered only in the form of full, off-the-shelf, or customized courseware. However, users also need a way to efficiently turn their proprietary knowledge into effective eLearning content through content-authoring

FIGURE 4

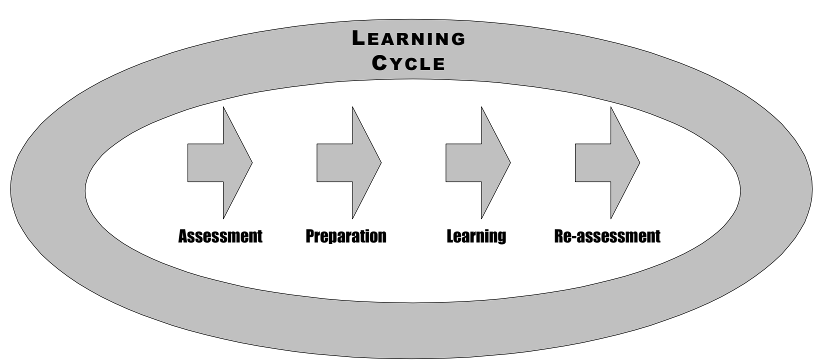

tools. While general knowledge provides a necessary foundation, proprietary knowledge provides organizations with competitive advantage.

Learning Cycle

As work, life, and learning continue to converge, portals will be the primary tool used to facilitate access to knowledge and learning. The ideal eLearning portal model supports the learning cycle from assessing needs through the learning process to learning evaluation. The learning cycle consists of four elements: assessment, preparation, learning, and re-assessment.

Full-Service eLearning Portal

The complete or full-service eLearning portal supports the learning cycle with different components of eLearning. Many of these components are foundational to the learning process and are critical in creating a full-service eLearning portal. All components of a full-service eLearning portal are fully integrated with seamless transition from one component to the next.

The full-service eLearning portal is comprised of three stages: assessment, competency, and learning evaluation. The assessment phase is composed of components for knowledge assessment, competency assessment, and learning evaluation. The preparation stage contains learning catalog, e-commerce, and enrollment components. The learning phase is comprised of learning activity, expert forum, and community components.

- *Assessment:* Assessing the learning needs of a user begins with an evaluation of his knowledge or competencies. This knowledge assessment is then compared with the competencies required for the job.
- *Preparation:* The user makes preparations for fulfilling his learning need through creating a plan from the list of learning activities that would best meet his need.
- *Learning:* The user engages in learning activities to build knowledge and develop competency (from expert forum to community collaboration.)

The learning-cycle revolution is accomplished when increased competency is verified through user evaluation. Underpinning a successful, full-service learning portal is the inclusion of an LMS.

ReLearning

eLearning combines traditional classroom instruction with technology-based delivery, offering blended learning solutions to the masses. eLearning offers the following:

- *Real-Time Learning:* Immediate and up-to-date application of critical knowledge and information.
- *Learner-Centric Training:* Training focus changes from traditional instructor-centric to learner-centric training. Instruction is tailored to the users professional responsibilities and capabilities.
- *Attract, Train, and Retain:* Addresses the need for the user to develop new knowledge and skills by providing the user with learning on demand.
- *Personalized Individual Training:* System is adaptive; it learns about its users and tailors its offerings to their learning style, job requirements, career goals, current knowledge, and personal preferences.
- *Ownership and Empowerment:* Users are responsible for their own learning—empowering them to manage and implement their own learning and development plans.
- *Simulation:* Introduces to the user the required, "interactive" part of learning—interaction and participation with a local or global audience.
- *Collaboration:* Accomplished through either joint problem solving or discussion among study groups through forums, discussion groups, and chat rooms—opening the path to broader thought and innovative processes through the sharing of ideas and experience.
- *Anytime, Anywhere:* The reality of training in a virtual information classroom across continents is now possible—anytime, anywhere. eLearning is less intrusive to the user, saving time and money through less interruption to the user.
- *Cost-Effective:* Costs can be measured and applied to each user and results can be measured against the incurred costs.
- *Metrics:* eLearning return on investment (ROI) can be effectively measured in terms of knowledge gain and retention.

Portal Technology

Portal technologies supporting open standards can be easily integrated into an organization's existing infrastructure. The portal needs to be operating system and Web-server–neutral so that enterprises can host it on the platform of choice. The portal solution should be deployable and accessible across a variety of platforms and devices. With a platform-, application-, and device-independent architecture, the eLearning portal provides optimal flexibility.

A modular approach provides the greatest flexibility and efficiency for building content, collaboration, and commerce functionality. Moreover, the current business climate and economies demand that enterprise technology have the capability to adapt to changes in the user base and integrate with the most demanding applications. Many eLearning portals have been built from the ground up to be a true enterprise-strength solution. This allows organizations to implement with confidence, knowing that their portal server can accommodate not only thousands, but also hundreds of thousands of users if required.

Next-Generation Portals

The ability to provide mobile, distributed workers with the organized access to the applications, knowledge, and information they need for sound decision-making has become vitally important for businesses striving to be productive, agile, and profitable. The attractiveness of Web-based computing, combined with the need to expedite information access and learning, has fueled the adoption of eLearning portals.

Open technology architecture will enable application access on virtually any device, including wireless and handheld communication devices and information appliances.

Mobile users will be able to move seamlessly from one device to another and receive consistent, personalized learning and knowledge.

Future eLearning portal features coming to market include better process integration, cascading portals, federated portals, business intelligence, and knowledge management. eLearning portals will connect directly and seamlessly with ERP, business intelligence, CRM, and other mission-critical enterprise systems.

Conclusion

The Internet and the Web have marshaled in an unprecedented business and knowledge revolution. A revolution that represents a fundamental shift in the way business is conducted and managed. Over the last quarter century or so, the industrial world has transitioned from being deprived of data to being besieged by it.

With the Internet came speed, connectivity, and intangible value—and the ability to "e-enable" all facets of business, including learning, knowledge, and performance management. But with the e-enabled organization comes stockpiles of data and information—overwhelming the organization and learner.

eLearning portals will drive the evolution from the information economy to the knowledge economy. Moreover, new technologies and the power and connectivity of the Internet will enable eLearning technology, content, and service companies to develop critical learning resources—revolutionizing the way we mentor, train, and learn. eLearning portals will be portable. The learning portal will be positioned as an integration and development platform, not a separate standalone application. The continued knowledge revolution will allow the eLearning portal to bring all information into a distinct, consistent, easily used interface while being fully integrated with other enterprise systems.

Submarine Applications Utilizing Ultra-Long-Reach Transmission Technology

Paul Bullock
Manager, Next-Generation Architecture
Nortel Networks

Niall Robinson
Director, Product and Technology Planning
Nortel Networks

Siraj Nour El-Ahmadi
Vice President, Business Development and Product Planning
Nortel Networks

Introduction

With the recent advancements of terrestrial ultra-long-reach transmission systems, new applications in a relatively untapped market for festoon and regional undersea networks are rapidly emerging. Festooning, a technique for installing unrepeatered submarine systems along a coastline, may be warranted in cases when terrestrial routes are physically impassable, uneconomical, or politically unstable, thus making right-of-way access unattainable. To meet the requirements of service providers operating backbone networks over large geographic areas, the ability to seamlessly complete networks utilizing both terrestrial and submarine technology is becoming critical. To achieve that goal, transmission systems must effectively be able to transport high capacities of traffic all-optically among strategic backbone hubs, sometimes separated by thousands of kilometers. In many cases, the head-end terminations along a festoon route provide excellent access points for transoceanic traffic, thus establishing the single-hop undersea and festoon system as somewhat of a collector infrastructure for global optical networks.

Validated through extensive field experimentation and commercial deployment, ultra-long all-optical distances are achieved utilizing a combination of nonlinear dispersion managed return-to-zero (RZ) transmission, distributed Raman amplification, and strong out-of-band forward error correction coding (OOB–FEC) [1][2]. Borrowing from these same fundamental techniques, it is now possible to support a variety of unrepeatered single-hop festoon and regional undersea applications much more efficiently and cost-effectively than traditionally possible. Submarine adapted ultra-long-haul terrestrial-based systems are advantageous in that they eliminate the need for active equipment in the wet-plant portion of the system. Traditional repeatered submarine systems rely on highly customized undersea technology, including optical repeaters designed to operate continuously and without maintenance activities for as long as 25 years. Such prerequisites command high premiums for the handpicked, high-reliability components required for such applications. In addition, because of the stringent physical packaging and powering constraints of undersea repeaters, they present a bottleneck that keeps undersea transmission systems from reaching the bandwidth capacities of their terrestrial counterparts.

In summary, this paper examines many of the economic and technical considerations surrounding regional unrepeatered undersea applications and how recent advancements in terrestrial transmission adapt to support this emerging market segment.

Regional Undersea Applications

As carriers strive to offer global connectivity, submarine and terrestrial networks continue to converge on utilizing the same technologies. Just as synchronous digital hierarch (SDH) systems have replaced asynchronous equipment and ring topologies have replaced point-to-point (PTP) systems, so too has the recent technological advancements in terrestrial ultra long haul made their way into unrepeatered undersea applications [3]. While submarine carriers have utilized techniques such as OOB–FEC and uniform low-power erbium-doped fiber amplifier (EDFA) deployment to increase system reach, the commercialization of Raman-assisted optical amplification and RZ modulated transmission will further increase the capabilities of such systems. Unrepeatered submarine systems are very similar to terrestrial systems in that all active components are located on dry-land locations for ease of accessibility and routine maintenance. Whenever the interconnecting sites are within reachable proximity, such single-hop systems are always economically preferable compared to traditional repeatered submarine systems. With current technology, it is possible to transport multiple 10 gigabit per second (Gbps) dense wavelength division multiplexing (DWDM) circuits over distances of more than 300 km without the use of remote repeaters.

There are a number of economic and operational benefits gained by utilizing unrepeatered undersea systems over more traditional repeatered undersea systems. Commercial applications for such systems include channel and sea crossings in Europe, island-to-island crossings in the Asia-Pacific and the Caribbean, as well as coastal festoon systems in South America, North America, and elsewhere in the world. With the emergence of commercial ultra-long-haul technology, the same platforms can support both land and submarine applications. Land-based systems utilize the reach capability to eliminate remote electronic regeneration, while submarine systems benefit from the elimination of undersea amplification elements. Both systems are monitored under the same management platform. As shown in *Figure 1*, carriers may utilize unrepeatered systems in conjunction with terrestrial systems to offer route diversity under the same management platform. A coastal festoon hops along a coastline and terminates periodically to accommodate reach constraints or drop traffic. Also, a single-span undersea system is shown bridging the mainland facility with a remote island.

Unlike repeatered undersea systems, each physical cable is actually capable of accommodating hundreds of fiber strands. However, with traditional repeatered systems, the number of active repeater modules that can fit within the undersea housing becomes the limiting factor in overall system capacity. Today, repeatered systems are limited to eight fiber pairs, while unrepeatered systems can accommodate more than 100 fibers, thus allowing for significantly more capacity. Submarine repeaters are made of optical amplifiers packaged in undersea-resistant housings made to withstand sea pressure and other stresses related to laying and recovery. Within these repeaters reside specialized optical amplifiers designed for maximum reliability and an in-service life span of as long as 25 years. Such strict engineering tolerances and customized requirements can drive the price of each undersea amplifier toward U.S.$1 million. The inclusion of Raman pumps in these repeatered undersea systems may help to increase the number of wavelengths supported on each fiber but will also further constrain the overall number of fibers each amplifier housing can accommodate while reducing the system availability by imposing additional active elements into the system. In addition to the physical housing constraints of undersea repeaters, specialized power systems also create a bottleneck for capacity scalability. Active repeaters require specialized power-feed equipment, which provides high-voltage direct current (DC) (as much as 10,000 volts) and regulated current (as much as 1.6 amps). DC power availability may be the limiting factor in

FIGURE 1

Utilizing commercial ultra-long-haul technology carriers can support both land and submarine applications utilizing a common management and operations platform.

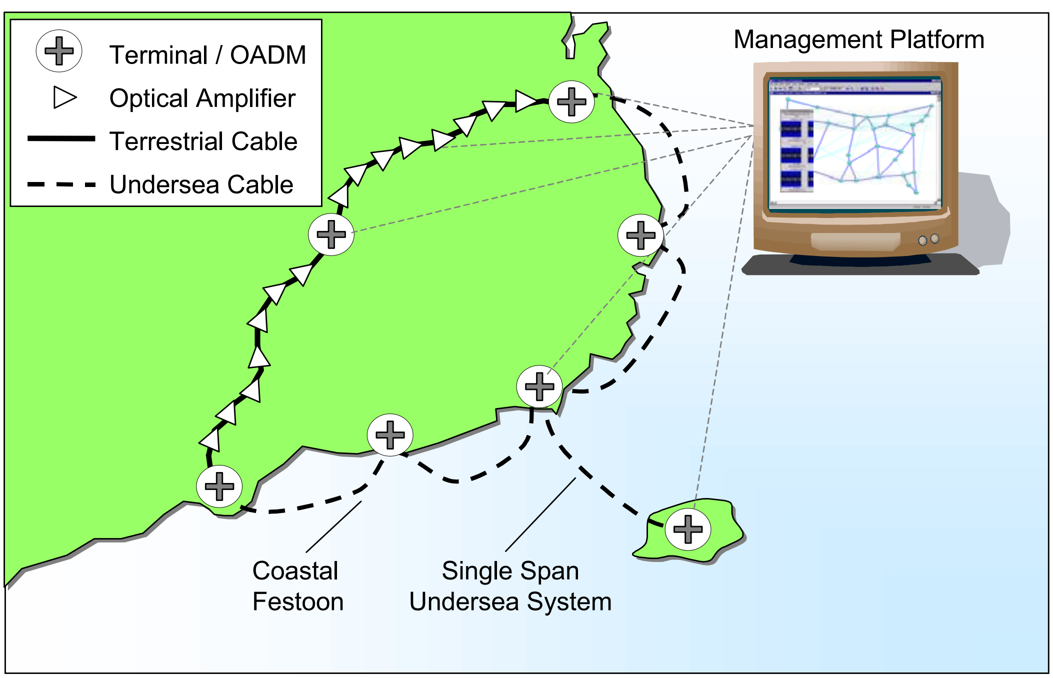

the race for ever-increasing capacity. Features include redundancy to withstand component failures or for routine maintenance, safety features to prevent accidental access to high-voltage circuits, and maintenance features to switch power feeds and low-frequency current-modulation capabilities to allow cable-ship repair teams to locate buried cable [4]. With unrepeatered systems, there is no need to take into consideration the physical and reliability constraints of repeaters such as customized undersea housing and strict optical amplifier tolerances. Thus, the elimination of undersea amplification in regional submarine applications dramatically reduces operational costs while significantly improving deployment velocity, thereby enabling carriers to quickly realize revenue.

In addition, by utilizing the extensive optical add/drop functionality of advanced transmission platforms, it is possible, depending on the distance of the undersea span, to employ these capabilities at terrestrial landing sites to further extend the reach of the system and optically integrate with the terrestrial portion of the network. Advancing beyond the first generation of optical add/drop-capable transmission systems, today's platforms, based on the enabling technologies described previously, support fully flexible optical add/drop capability [5]. All-optical branching allows for the re-direction of wavelengths at the landing point to multiple fiber directions without introducing the need for electrical termination. Such capabilities enable carriers with the potential to optically re-direct wavelengths as they glass-through the terrestrial landing site and continue all-optical transmission to the traffic destination sites, which may be several hundred kilometers away. Such flexibility not only reduces electrical termination at the landing site, thus reducing capital expenditures (CAPEX) and operational cost, but also enables rapid service deployment.

Global carriers are finding a number of advantages to utilizing common technologies and platforms for both the terrestrial and submarine portions of their network to simplify operations and management. Fully integrating the regional undersea and terrestrial transmission requirements under a common platform enables simplified training and maintenance procedures along with streamlined network management and system monitoring. Furthermore, without the need for specialized active undersea components, submarine adapted ultra-long-haul systems improve delivery and in-service time frames, which clearly translates into a faster return on investment (ROI) for new projects. *Table 1* summarizes the advantages of extended-reach single-hop systems compared to traditional repeatered submarine systems for regional undersea applications.

Technical Considerations

Unrepeatered systems are typically limited in distance by the need to balance the effects of pulse distortion through nonlinearity induced by high signal powers with optical signal-to-noise ratio limitations of the received signal. Other impairments, such as polarization mode dispersion, polarization dependent gain, and amplifier gain tilt can generally be ignored due to the short length of these systems in comparison to ultra-long-haul terrestrial and repeatered undersea systems.

TABLE 1

Comparison of unrepeatered extended-reach single-span systems and traditional repeatered systems for regional undersea applications.

	Extended-Reach Single-Span Systems	*Regional Repeatered Submarine Systems*
Utilizes active undersea amplifiers?	No	Yes
Customized components required for operation?	No, same as terrestrial system	Yes, route-specific and designed to operate untouched
Specialized power systems?	No, same as terrestrial system	Yes, high-voltage DC with expensive redundancy
Specialized operations?	No, same as terrestrial system	Yes, more difficult to manage, maintain, and repair
Cable-capacity limitation?	As many as 100 fibers	Currently eight fibers
Seamlessly integrates with terrestrial systems?	Yes, optically compatible with terrestrial systems	No, often different systems used for terrestrial and wet plant
Shares terrestrial management platform?	Yes, common management platform and monitoring	No, undersea systems are isolated from terrestrial systems
All-optical add/drop capability at landing sites?	Is possible, depending on reach, to optically branch wavelengths for terrestrial transmission	No, full electrical termination required at landing sites

However, transmission techniques used to extend terrestrial and repeatered submarine systems have a role to play in improving the capacity-reach product of unrepeatered submarine links. At the launch end of the transmission line, self-phase modulation is generally the limiting nonlinearity. However, as capacity levels increase and therefore the total optical bandwidth increases, the signal-to-signal power transfer through the Raman gain process in the transmission fiber will begin to be the limiting factor. Signal-power-level pre-emphasis techniques used to optimize ultra-long-haul systems will be required to combat this effect.

An interesting development of unrepeatered systems is the ability to use different fiber types in the transmission link. The nonlinearity level induced in a fiber is, to an order, determined by the effective area of the optical fiber and the signal-power level. The effective area of an optical fiber is a representation of the cross-section of the optical fiber that is actually guiding the transmitted light. Fibers with small effective areas squeeze the light into a smaller cross-section, therefore increasing the intensity of the light creating a higher nonlinear effect for the same power level compared to a fiber with a larger effective area. These effective-area managed cables [6] often incorporate chromatic dispersion management into the design. An example of an effective-area managed cable might be to use a fiber with a very large effective area at the transmission end of the system to increase the allowable launch power level before effects such as self-phase modulation and channel-to-channel Raman gain become the limiting factor. At the receive end of the link, a section of fiber similar to traditional dispersion-compensating fiber could be used. The advantage here is the combination of built-in dispersion compensation in the transmission fiber and a small effective area that enables large distributed Raman gain to be achieved at relatively low levels of Raman pump power.

Although not new to the world of submarine systems, OOB–FEC techniques can provide several decibels (dBs) of additional margin, some of which could be used to extend the overall system reach. Early synchronous optical network (SONET)/SDH–based error-correction implementations relied on the utilization of normally unused overhead bytes in the transmission frame. Such "in-band" techniques are able to provide 2 to 3 dB of net error-coding gain. However, by manipulating overhead information, such traditional FEC methods are unsuitable for the completely transparent transport of client signals, which may also utilize the same SONET/SDH overhead for proprietary control or messaging information. Current submarine and some emerging terrestrial systems use "out-of-band" FEC at line rates of up to 10.7 Gbps. Such, coding schemes yield a 4 to 5dB net error-coding gain, which is certainly an improvement over the "in-band" technique. Such methods have been commercially utilized in repeatered undersea systems. The most advanced commercially available ultra-long-haul transmission platform increases the normal 9.96 Gbps line rate to 12.24 Gbps to deliver a best-in-class 9 dB net error-coding gain at a 10^{-15} bit-error rate. This out-of-band implementation allows for the completely transparent transmission of client traffic, while also providing critical performance-monitoring capabilities. Such strong OOB-FEC enables coding gains that are higher than those found in repeatered undersea systems. OOB-FEC encapsulation is well suited for a variety of transmission formats such as SONET, SDH, and 10 Gbps Ethernet. The OOB–FEC overhead can also support the optical service channel (OSC) information that would normally reside in an out-of-band wavelength alongside the C-Band window of operation for terrestrial transmission.

In addition to strong FEC, terrestrial transmission systems have commercialized the use of Raman amplification, which has also found its way into the unrepeatered submarine market. With single-span unrepeatered systems, the Raman pump laser is located at the far end of the system and is pumped into the transmission fiber counter—spreading to the traffic-carrying wavelengths to provide a distributed gain along the fiber medium itself. The Raman laser excites the transmission fiber, thus effectively boosting the transiting signal level as it travels down the fiber closer toward the Raman pump source. Here, the ultimate limitation will be determined by the fiber fuse limit of the transmission cable. Fiber fusing is a runaway effect, where the optical power level launched into the fiber is sufficiently high to cause permanent physical damage to the transmission fiber rendering it unusable for transmission.

Many ultra-long-haul transmission systems, both terrestrial and submarine, utilize some form of a RZ modulation format. In unrepeatered links, the advantage of traditional RZ modulation is not as evident. Typically, RZ transmission requires a lower optical signal-to-noise ratio (OSNR) at the receiver in comparison to non–RZ (NRZ) transmission. However, with the advancement in potential modulation techniques now available in optical transmission systems such as carrier-suppressed RZ and vestigial sideband RZ, further improvements in system reach or increases in the transmitted line rate may be possible.

As shown in *Figure 2*, the unrepeatered extended reach single-span system encapsulates the original 10 Gbps-client signal in the payload of the out-of-band FEC frame. Error-correcting overhead codes for the client signal are appended, and the entire frame is externally modulated in an RZ format and launched onto the fiber along with a number of other wavelengths. Following propagation through the fiber medium itself, the incoming wavelengths are gradually amplified by the counter-pumped Raman source and discretely boosted by the post EDFA. Each wavelength is demultiplexed, the client signal is decoupled from the error-correction frame, and any errors detected during transmission are monitored and corrected.

The single-span system described could be augmented with remote optically pumped amplifiers (ROPAs) to further improve the reach and capacity of the transmission system. ROPAs are optical repeaters without active elements contained in the underwater housing. A single coil of erbium-doped fiber and associated isolators are housed in a submarine qualified repeater housing spliced into the optical cable. In a ROPA configuration, terminal generated optical pump power is transmitted to the erbium gain block located 80 to 100 km away from the terminal site along the fiber cable so there are no electrically active components in the fiber plant, thus a high-voltage DC power feed is not required. Although ROPA systems are not strictly "repeaterless," because optical signal amplification occurs in the remote gain block, they are not as restrictive as traditional "repeatered" undersea systems. ROPA systems, however,

FIGURE 2

Unrepeatered regional submarine system architecture that enables the transparent transmission of 10 Gbps client signals utilizing OOB–FEC, dispersion managed RZ modulation, and a combination of EDFA/Raman amplification. Only one direction is shown for clarity.

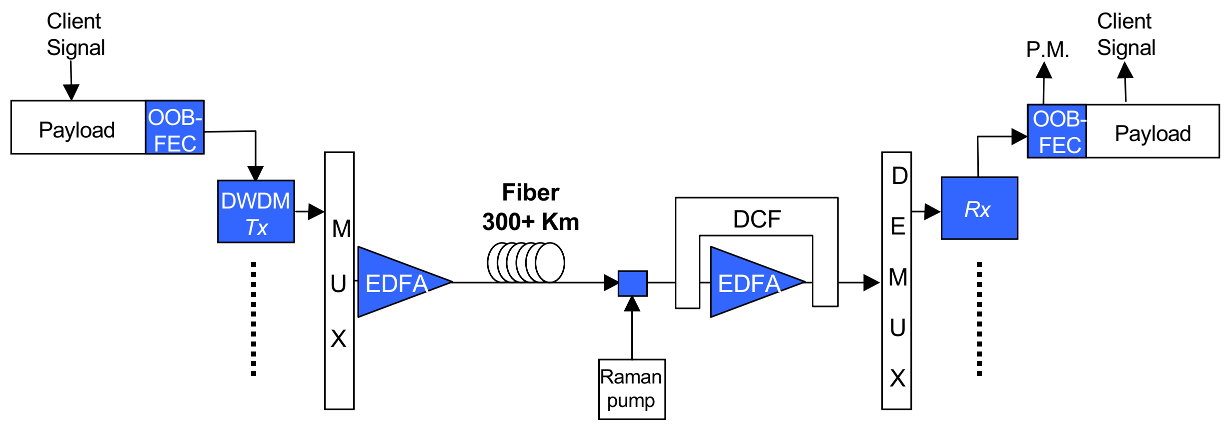

still require specialized physical considerations to withstand undersea conditions and restrictive reliability constraints and thus can create a significant price premium over unrepeatered systems [7].

Market Outlook

Although several coastal festoon and regional undersea systems have been in place for many years, these systems have quite often relied on the use of specialized undersea repeaters, thus they were limited by the packaging and power constraints of active undersea components. With the recent advancements in ultra-long-haul transmission systems, it is now possible for carriers to cost-effectively support these applications, along with many other opportunities. As the reach, capacity, and manageability of unrepeatered systems continues to scale beyond traditional repeatered deployments, the market outlook continues to broaden, as previously cost-prohibitive ventures throughout the world have now become viable business models. Furthermore, by linking with terrestrial systems through homogeneous platforms, a much broader range of applications and potential for long-term growth emerges. According to recent market forecasts, the global market for unrepeatered transmission will grow substantially from 2001 through 2005. More than 300 unrepeatered links are installed or planned throughout the globe, with some 120 additional links targeted for future deployment, scaling the market from a current $2.5 billion to more than $7 billion by 2005 [8].

Conclusion

The global market for festoon and regional undersea systems is new and relatively untapped. Utilizing commercially available ultra-long-reach, terrestrial technology carriers are now able to rapidly and cost-effectively establish presence in a number of markets along and between coastlines. Enabled by strong error-correction coding, Raman amplification, and advanced modulations techniques, it is now possible to transparently support native 10 Gbps DWDM systems to maximize the bandwidth-distance product without incurring the operational and CAPEX of deploying and maintaining specialized undersea optical repeaters. Borrowing from such recently commercialized advancements in ultra-long-reach terrestrial transmission systems, unrepeatered applications can utilize the same techniques to mitigate nonlinear effects to enable native 10 Gbps unrepeatered transmission between landing sites separated by hundreds of kilometers. Seamless integration with the terrestrial plant ensures simplified operational and maintenance procedures, thus reducing the life-cycle cost of the system, while allowing for diverse optical protection switching schemes. This next generation of unrepeatered transmission systems can augment existing network capacity, provide alternate protection routes for land-based systems, or completely obsolete many of the systems currently in use by supporting higher capacities in a more cost-effective manner than previously possible.

References

[1] Haxell, I., et al., "2410Km All-Optical Network Field Trial with 10 Gb/s DWDM Transmission," Optical Fiber Communication Conference (OFC '00), PD41 (Baltimore, Md., March 7–10, 2000).

[2] "Genuity Completes Industry's First All-Optical Network Segment." Genuity Press Release. 12 December, 2000. <http://www.genuity.com/announcements/news/press_release_20001204-02.xml>.

[3] Teral, S. R., "Global Optical Networking with Undersea and Terrestrial Networks," National Fiber Optic Engineers Conference (NFOEC '01) (Denver, Colo., August 2000).

[4] Welt, D., Kostic, A., Fadlallah, I., "Submarine Cable Systems Tutorial," Nortel Networks (January 2001).

[5] Jordan, K., Bullock, P., "Extended Reach Optical Add-Drop Applications and Drivers."

Optical Transport Network Recovery through Photonic Switching and GMPLS

Luc Ceuppens
Vice President, Corporate Business Development
Calient Networks

Optical amplification and dense wavelength division multiplexing (DWDM) have fundamentally changed the economics of optical transport networks. Now that these technologies are widely adopted, the bottleneck has moved from the outside line plant to nodal central offices (COs), where the speed of electrical switching equipment has not kept pace. As a result, present networks are carrying terabit per second (Tbps) data streams as 100 parallel 10 gigabit per second (Gbps) or 500 parallel 2 Gbps streams. Each stream is terminated within the CO and carried over individual strands of fiber, creating operational and financial nightmares. This is primarily due to the cost of acquiring and operating the countless optical-to-electrical-to-optical (O–E–O) devices in the signal path inside the CO. While O–E–O technology was (and still is) necessary for the grooming and aggregation of low-speed services onto wavelengths, the transport network has dramatically changed. The inflow of data traffic is threatening to saturate the existing networks, and this requires a dramatic rethinking of how networks need to be planned, designed, constructed, and operated. While today's transport networks carry remarkable amounts of bandwidth, the optical layer of these networks is fundamentally static and provides for only simple point-to-point transport. While efficiently addressing the capacity problem, DWDM alone does little to make the networks smarter. Efficiently managing (i.e., adding, dropping, routing, protecting, and restoring) the growing number of traffic-bearing wavelengths can only be achieved through a new breed of networking element: Photonic switching systems (PSSs) can efficiently execute these functions because they are bit-rate, wavelength, and data-protocol transparent. With their all-optical switch cores and interfaces, PSSs can switch and manage optical signals at various levels of granularity—wavelength, sub-band, full-band, and composite DWDM fiber levels. Though cross-connect systems with electrical switch cores [1] are available, they perform these functions at very high capital costs and operational inefficiencies. This paper examines enabling technologies for the deployment of PSSs in carrier optical transport networks (OTNs) and takes a practical perspective on survivability architecture migration and implementation issues.

The Challenges of Building, Operating, and Managing Optical Transport Networks

For much of the late 1990s and through the early part of 2000, competitive local-exchange carriers (CLECs), wholesale broadband providers, and high-speed Internet companies raced to build new networks and add customers while demand for high-speed services soared. Then, during the second half of 2000, market conditions and a wary investor community caused the large incumbent carriers to restructure their businesses, while many smaller competitors fell by the wayside as capital markets dried up and profits became more important than growth. The telecommunications downturn was precipitated by the suffering stock markets, which led to a shortfall in investment capital as Wall Street and venture capitalists became more restless. Suddenly, instead of focusing on speedy growth, companies were scrambling for enough capital to survive. Those still standing as 2002 arrives are faced with shaking free from the recent carnage, making their new business focuses successful, and renewing investor confidence. The primary trigger for this will be better operating results.

Building and maintaining a network infrastructure is a costly business, and with the markets remaining hostile to new borrowing, communications companies continue to face hardships. Nonetheless, service providers cannot afford not to grow their networks. So, while capital spending is slowing down, the money that is available will have to be spent intelligently and carriers will have to focus on faster, cheaper, and smarter systems. While growing the network, every dollar invested will have to dramatically reduce the operating cost of both new and existing network segments. At the same time, the new infrastructure has to support new revenue opportunities. This leaves service providers with three challenges:

- Reduce and optimize capital expenditures
- Reduce operating expenditures
- Identify and execute on new revenue opportunities

In today's competitive business environment, any technical strategy for network evolution must ultimately provide for competitive differentiation in service delivery. Scaling the network and delivering bandwidth and services when and where a customer needs it are absolute prerequisites for success. The limitations of the existing network infrastructure are hindering a movement to this service-delivery business model. A new network foundation is required—one that will easily adapt to support rapid growth, change, and a highly responsive service delivery.

Photonic switches offer efficient solutions and improvements by providing intelligent optical bandwidth management in support of fluid bandwidth creation and distribution. They provide numerous financial, operational, and tactical benefits by switching signals at the optical level and by providing efficient optical bandwidth management to enable new services and gain operational efficiencies.

Enabling Technologies

Two enabling technologies make large-scale PSSs commercially deployable: 1) the availability of large three dimensional (or 3D) micro-electromechanical system (MEMS) photonic switch components and 2) nearly complete standardization of control plane protocols, which is key to the implementation of PSSs in service-provider networks.

MEMS Photonic Switches

Since 1980, various photonic switching technologies have been developed, including integrated optic, thermo-optic, acoustic-optic, total internal reflection, and liquid-crystal–based switches. These technologies have suffered problems, such as high insertion loss, wavelength and temperature dependency, small scale, and polarization dependency [2], which hamper their adoption in telecommunication networks. 3D MEMS mirror technology is the only photonic switch technology to offer the scalability, reliability, non-blocking operation, fast switching, low insertion loss, low crosstalk, and high serviceability features that service providers require. These switches (shown in *Figure 1*) offer bit-rate, protocol, and wavelength independence between 1260 and 1625 nm. In these large-scale free space switches, light beams are switched with mirror arrays that pivot about two orthogonal axes [3]. As laser beams are spatially and temporally coherent, they do not interact when crossed, hence blocking does not occur. This technology also offers polarization independency with very low polarization mode dispersion and polarization dependent loss.

Generalized Multiprotocol Label Switching

Network protection, dynamic restoration, and real-time dynamic service provisioning of wavelengths can be implemented using generalized multiprotocol label switching (GMPLS), which combines existing control-plane techniques with the point-and-click provisioning capabilities of photonic switches [4]. GMPLS is a signaling and routing protocol suite, with extensions to some existing Internet protocol (IP) protocols. GMPLS is a logical extension of multiprotocol label switching (MPLS), the set of extensions to open shortest path first (OSPF), intermediate system to intermediate system (IS–IS), and resource reservation protocol (RSVP) to support the routing of paths (also known as traffic engineering). Intrinsic to MPLS is the notion of separating the forwarding (data) plane from the control (signaling) plane. Data is transported across the network inside a path. Multiprotocol lambda switching (MPλS) was a concept that said that the MPLS control plane could be leveraged to support routing of lambda paths. After all, there are several synergies between label-switched routers (LSRs)

FIGURE 1

Switching in a 3D MEMS Photonic Switch

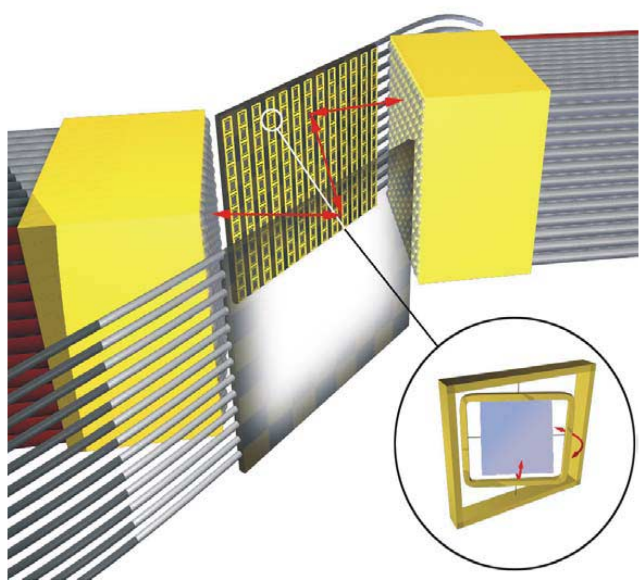

and photonic switches, and between a label-switched paths (LSP) and an optical trail. Analogous to switching labels in an LSR, a photonic switch switches wavelengths from an input to an output port. Establishing an LSP involves configuring each intermediate LSR to map a particular input label and port to an output label and port. Similarly, the process of establishing an optical path involves configuring each intermediate photonic switch to map a particular input lambda and port to an output lambda and port. As in LSRs, photonic switches need routing protocols (such as OSPF and IS–IS) to exchange link-state topology and other optical resource availability information for path computation. They also need signaling protocols to automate the path establishment process.

GMPLS is the realization of the MPλS concept, created by extending MPLS to support non-packet paths. GMPLS offers generalized semantics for entities such as packets, timeslots, wavelengths or lambdas, wavebands, fibers, cables, and ducts. A critical component of GMPLS is the link management protocol (LMP), a new protocol that runs between adjacent nodes for both link provisioning and fault isolation [5]. The GMPLS protocol suite enables client-layer data network elements—such as routers, ATM switches, and synchronous digital hierarchy (SDH)/synchronous optical network (SONET) add/drop multiplexers (ADMs)—to dynamically request bandwidth services from the server photonic network layer (consisting of optical network elements such as PSSs and DWDM systems) by using existing control-plane techniques. Beyond simple advertising and the establishment of optical paths through the network, GMPLS includes performance information about the network. Client-layer equipment at the edges of the photonic network uses this information to dynamically request resources (for example, wavelengths) available in the server photonic layer. In essence, GMPLS reduces the network view to just two layers: the OTN server layer and the client service-delivery layer.

GMPLS–Based Network Recovery

GMPLS is a multipurpose control-plane paradigm that supports not only devices that perform packet switching, but also devices that perform switching in the time, wavelength, and space domains. To leverage the natural connection and device hierarchy that exists in real-world networks, GMPLS introduces the concept of hierarchical LSPs [6] (see *Figure 2*), which occurs when a new LSP is tunneled inside an existing higher-order LSP so that the pre-existing LSP serves as a link along the path of the new LSP. The ordering of LSPs is based on the link multiplexing capabilities [7] of the nodes. Nodes at the border of two regions, with respect to multiplexing capabilities, are responsible for forming higher-order LSPs and aggregating lower-order LSPs. As an example, the bandwidth signaled for LSP_1 in *Figure 2* is 500 megabits per second (Mbps); all links traversed by the LSP must be large enough to support the requested bandwidth. R_0 classifies

FIGURE 2

GMPLS LSP Hierarchy

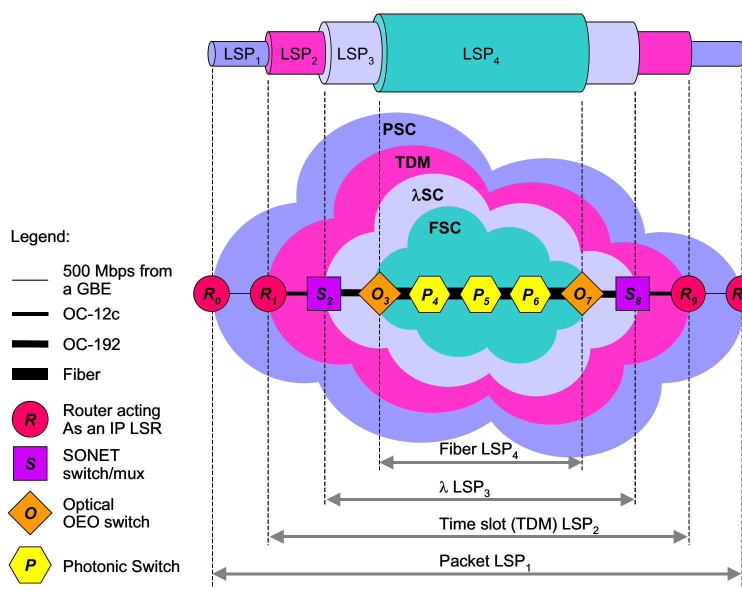

LSP_1 is set up from R_0 to R_{10} for 500 Mbps of bandwidth. LSP_1 is nested in the subordinate LSPs 2, 3, and 4, respectively, as depicted at the top of the figure. If an existing LSP has capacity, the subordinate LSP is added to it; otherwise, the set-up of a new LSP of the appropriate type is triggered. The endpoints of LSP_1 classify packet data and source it into LSP_1. R_1 and R_9 are packet LSRs; S_2 and S_8 are SONET path-level switches. Devices O_3 and O_7 are optical or lambda switched, providing SONET/SDH section-level signals (e.g., OC–192 including all overheads); they also have WDM capabilities between photonic switches P_4 and P_6. The link between R_0 and R_1 is Gigabit Ethernet, between R_1 and S_2 an OC–48, between S_2 and O_3 an OC–192, between O_3 and P_4 a WDM multiplex of 16 OC–192 that which remains intact through to O_7 (P_4 to P_6 are pure photonic switches). The link between O_7 and S_8 is an OC–192, between S_8 and R_9 an OC–48, and between R_9 and R_{10} Gigabit Ethernet.

and maps packets into LSP_1 and shapes the data to 500 Mbps if necessary. The 500 Mbps for LSP_1 will be allocated from LSP_2, which is running at the next larger SONET increment, an optical carrier (OC)–12c. The S_2 switch grooms the OC–12c into LSP_3, which is an OC–192 between S_2 and O_3. The optical switch, O_3, takes this OC–192 and switches it through LSP_4, a WDM channel toward P_4. The OC–192 corresponding to LSP_3 is switched through P_4, P_5, and P_6 to O_7. Continuing in this fashion, O_7 selects the correct lambda and passes the signal to the port adjacent to S_8. S_8 will select the appropriate OC–12c from the OC–192 and pass this to R_9. Finally, R_9 will take the packets from the OC–12c and forward them on R_{10}.

The GMPLS protocol suite broadly encompasses routing (OSPF–TE/IS–IS–TE), signaling (RSVP–TE/CR–LDP), and link management (LMP). These protocols interact in a meaningful manner to provide coordinated hierarchical recovery of connections across multiple layers, something that cannot be achieved if restoration is driven by each layer individually. GMPLS–based recovery provides an answer:

- When Layer-3 rerouting may be too slow
- When Layer-0 or Layer-1 mechanisms may have no visibility into higher-layer operations
- When lower layers may provide some levels of protection (for example, link protection) but no protection against node failures or compute disjoint paths
- To ensure interoperability of protection mechanisms between network elements from different vendors
- To provide a recovery option without an intervening SONET layer when IP traffic is transported directly over WDM optical channels

A key requirement for the development of a common control plane for both optical and electronic networks is the need for features in the signaling, routing, and link management protocols to enable intelligent fault management. Fault management requires four steps: fault detection, fault localization, fault notification, and fault recovery. Fault detection should be handled at the layer closest to the failure; for optical networks, this is the physical (optical) layer. One measure of fault detection at the physical layer is detecting loss of light (LOL). Other techniques based on, for example, optical signal-to-noise ratio (OSNR), bit-error rate (BER), dispersion, crosstalk, and attenuation are still being investigated (see, for example, [8] and [9]). Fault localization requires communication between nodes to determine where the failure has occurred (for example, SONET AIS is used to localize failures between SONET terminating devices). One interesting consequence of using LOL to detect failures in optical networks is that LOL propagates downstream along the connection's path. The LMP [10] includes a fault-localization procedure that is designed to localize failures in both transparent (all-optical) and opaque (opto-electrical) networks and is independent of the data-encoding scheme. Fault notification is the communication of a failure between the node detecting it and a node equipped to deal with the failure. Fast fault notification is essential for rapid recovery. The Notify mechanism [11] is designed to support fast notification of non-adjacent nodes.

Once a failure has been detected and localized, and the responsible node has been notified, protection and restoration can be used to recover from the failure. We make the distinction between protection and restoration by the time scales in which they operate. Protection is designed to react to failures rapidly (say, in less than a few hundred milliseconds) and often involves 100 percent resource redundancy. For example, SONET automatic protection switching (APS) is designed to switch the traffic from a primary (working) path to a secondary (protection) path in less than 50 ms. This requires simultaneous transmission along both the primary and secondary paths (called 1+1 protection) with a selector at the receiving node and uses twice as many network resources as a non–APS protected path. Restoration, on the other hand, is designed to react to failures quickly, but it typically takes an order of magnitude longer to restore the connection as compared to protection switching. This is because restoration typically utilizes pools of shared resources to improve network utilization efficiency. In addition, restoration may involve rerouting connections, which can be computationally expensive if the paths are not precalculated or if the precalculated resources are no longer available.

Protection and restoration methods have traditionally been addressed using two techniques: 1) path-level recovery, where the failure is addressed at the end nodes (i.e., the initiating and terminating nodes of the path), and 2) span-level recovery, where the failure is addressed at an intermediate or transit node. Path-level recovery can be further subdivided into path protection, where secondary (or protection) paths are pre-allocated, and path restoration, where connections are rerouted, either dynamically or using precalculated (but not pre-allocated) paths. Span-level recovery can be subdivided into span protection, where traffic is switched to an alternate channel or link connecting the same two nodes, and span restoration, where traffic is switched to an alternate route between the two nodes (this involves passing through additional intermediate nodes).

To effectively use protection, there must be mechanisms to configure protected links on a span between nodes, advertise the protection bandwidth of a link so that it may be used by a class of traffic that has different availability requirements, establish secondary (protection) LSPs to protect primary LSPs, allow the resources of secondary LSPs to be used by lower-priority traffic until a switchover occurs, and signal protection switchover when necessary. In the remainder of this document, we discuss protection and restoration in the context of GMPLS signaling.

Protection Mechanisms

Protection is designed to react to failures in the fastest timescale and typically involves pre-provisioning protection resources. The level of protection available is a function of the protection resources available for protecting a failed resource:

- *1+1 Protection*
  Two pre-provisioned resources are used in parallel. For example, data is transmitted simultaneously on two parallel links and a selector is used at the receiving node to choose the best signal.

- *1:1 Protection*
  Two resources (1 primary, 1 back-up) are pre-provisioned. If the primary resource fails, the data is then switched to the back-up resource. The recovery path can be used only to recover a specific working path.

- *1:N Protection*
  In 1:N protection, up to N working paths are protected using only one recovery path. N+1 resources (N primary, 1 back-up) are pre-provisioned. If there is a failure on any one of the primary resources, the data is then switched to the back-up resource. At this point, the remaining N-1 primaries are no longer protected.

- *M:N Protection*
  In M:N protection, up to N working paths are protected using M recovery paths. N+M resources (N primary, M back-up) are pre-provisioned. If there is a failure on any one of the primary resources, the connection is then switched to the back-up resource.

  Note that 1:N and 1:1 are special cases of M:N protection.

- *Span Protection*
  A span consists of a number of channels between two adjacent nodes that are grouped together into a single link, often called a traffic engineering (TE) link (see [10]). Span protection involves switching to a protection channel when a failure occurs on a working channel. At the span level, both dedicated (1+1, 1:1) and shared (M:N) protection may be implemented. The protection type supported by a TE link (LPT) is advertised throughout the network so that intelligent routing decisions can be made. The desired protection for a path is signaled as part of the generalized label request in GMPLS signaling. This is needed in signaling if a link supports multiple protection types or if loose routing is used.

  For dedicated 1+1 span protection, each node must replicate the data onto two separate channels (possibly using separate component links of a bundled link or separate ports of a TE link), and the adjacent node must select the data from only one channel based on the signal integrity. This is the fastest protection mechanism; however, it requires using twice the LSP bandwidth between each pair of nodes and the ability to replicate the data on two separate channels.

  For shared M:N protection, M protection links are shared between N primary links. Since data is not replicated on both the primary and secondary links, failures must first be localized before the switchover can occur. LMP can be used for fault localization, and the upstream node (upstream in terms of the direction an RSVP path message traverses) will initiate the local span protection.

- *Path Protection*
  Path protection is addressed at the end nodes of an LSP (i.e., LSP initiator and terminator) and requires switching to an alternate path when a failure occurs. For 1+1 path protection, a signal is transmitted simultaneously over two disjoint paths and a selector is used at the receiving node to choose the better signal. For M:N path protection, N primary signals are transmitted along disjoint paths, and M secondary paths are pre-established for shared protection switching among the N primary paths.

  After the two paths are computed, the source originates two explicitly routed connections with the dedicated 1+1 and unprotected bits, respectively, set in the protection bit vector of the corresponding signaling set-up message. The set-up indicates that these two paths desire shared reservations. For 1+1 path protection, the connection is transmitted simultaneously over two disjoint paths, and a selector is used at the terminator node to choose the best signal. At each node where the two paths branch out, the node must replicate the data into both branches. At each node where the two paths merge, the node must select the data from one path based on the integrity of the signal. For M:N path protection, N different connections are transmitted along N disjoint paths, and M disjoint paths are pre-established for shared protection switching for the N primary paths. An interesting feature of GMPLS is that it allows preconfiguring back-up paths to protect primary paths. These back-up paths, called secondary paths, are used for fast switchover when the primary path fails. Although the resources for these back-up paths are pre-allocated, lower-priority traffic may use the resources with the caveat that the lower-priority traffic will be pre-empted if there is a failure on the primary path. Note that the precomputed back-up path cuts down on the restoration time in the event of a failure.

Restoration Mechanisms
Restoration is designed to react to failures quickly and use bandwidth efficiently, but typically involves dynamic resource establishment and may also require route calculation and therefore takes more time to switch to an alternate path than protection techniques. Restoration can be implemented at the initiator node or at an intermediate node once the responsible node has been notified. Failure notification can be done using the Notify procedures [12] or using the standard RSVP PathError messages.

- *Span Restoration*
  To support span restoration, where traffic is switched to an alternate route around a failure, a new LSP is established at an intermediate node that involves passing through additional intermediate nodes. Span restoration may be beneficial for LSPs that span multiple hops and/or large distances because the latency incurred for failure notification may be significantly reduced and only segments of the LSP are rerouted instead of the entire path.

  If the protected part of the LSP is a single span, then error detection is sufficient to trigger restoration. If, however, protection is required over a series of more than one span, a mechanism is required to notify the point of repair that an error has occurred and that restoration is required.

  The RSVP Notify Request object can be used by an intermediate node to request that it be the target of an RSVP Notify message. Span restoration may break TE requirements if a strict-hop route is defined for the connection. Furthermore, the constraints used for routing the connection must be forwarded so that an intermediate node doing span restoration is able to calculate an appropriate alternate route. This is similar to the problems when establishing/maintaining TE requirements that span multi-areas (see [13] for a proposed mechanism).

- *Local Repair*
  Local repair is a special case of span protection supported by the base RSVP–TE draft [14]. The node that detects the failure may make an alternate routing decision and attempt to resignal the LSP. This approach may be considered too slow since it could rely on convergence of the routing table at the repair node. However, if there is a close link between routing and path computation components, local repair may be equivalent to span protection.

- *Path Restoration*
  Path restoration switches traffic to an alternate route around a failure, where the new route is selected at the LSP initiator and may reuse intermediate nodes used by the original LSP, and it may include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route calculation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the source node may not have enough information to reroute the connection around the failure. The back-up route may be calculated on demand (that is, when the failure occurs) or may be precalculated and stored for use when the failure is reported. This offers faster restoration time. There is, however, a risk that the back-up route will become outdated through other changes in the network—this can be mitigated to some extent by periodic recalculation of idle back-up routes.

Restoration (span or path) will be initiated by the node that has isolated the failure or by the node that has received either an RSVP Notify message or an RSVP PathError message indicating that a failure has occurred. The new resources can be established in a make-before-break fashion, where the new LSP is set up before the old LSP is torn down, using the mechanisms of the LSP_Tunnel Session object (see [14]) and the Shared-Explicit reservation style. Both the new and old LSPs share resources at nodes common to both LSPs. The Tunnel end point addresses, Tunnel Id, Extended Tunnel Id, Tunnel sender address, and LSP Id are all used to uniquely identify both the old and new LSPs; this ensures that new resources are established without double counting resource requirements along common segments. Note that make-before-break is not used to avoid disruption to the data flow (this has already been broken by the failure that is being repaired) but is valuable to retain the resources allocated on the original primary path that will be reused by the new primary path.

Restoration Classes
When the LSP initiator detects that the LSP carrying the traffic trunk has failed, it performs the action specified by the resilience attribute of the trunk. The possible actions are as follows:

- *None*
  No explicit path protection. This means that in failure scenarios, alternate ways have to be found to reroute the traffic.

- *Fallback to Precomputed, Pre-Established Path*
  This option requires that a back-up LSP be pre-computed and pre-established. This provides the fastest path restoration at the expense of wasting the bandwidth reserved by the back-up LSP.

- *Establish a Path Based on the Precomputed Path, and Move the Trunk onto the New Path*
  This option requires that the head-end router precompute a back-up path for the traffic trunk. Since the path is established on demand, there is no need to maintain idle paths. The price for the bandwidth efficiency over the previous option is the additional delay incurred for establishing the new path.

- *Compute and Establish a New Path, and Move the Trunk onto the New Path*
  This option provides the slowest restoration but incurs no bandwidth or processing overhead.

The response time of path restoration is equal to the sum of the following delay components:

- The time it takes for the router closest to the failure to detect it.
- The time it takes for the head-end router to detect the failure (via either an explicit RSVP failure notification or a PathError message).
- The time it takes to compute and establish a new path. Note that if the fallback path is precomputed and/or pre-established, this delay component is partially or completely eliminated.

The currently proposed timing bounds for service restoration for different mechanisms are as follows:

- M:N path restoration with pre-established capacity: 100 – 250 ms
- M:N path restoration with precomputed capacity: 100 – 750 ms
- Path restoration with computation: 1 – 5 seconds
- Local restoration: ≤ 50 ms

Routing Enhancements
The GMPLS extensions to OSPF [15] and IS–IS [16] include the advertisement of the LPT. The LPT field is a bit vector that indicates the protection capabilities that are supported for the link. The LPT field may be configured with dedicated 1+1, dedicated 1:1, shared M:N, and enhanced protection, as well as unprotected. For a link that has dedicated 1+1 protection or is unprotected, this advertisement provides a complete description of the link capabilities and the usable bandwidth. However, a key argument for using dedicated 1:1 or shared M:N is the efficiency gained by reusing the protection bandwidth for lower-priority traffic when the bandwidth would otherwise be idle.

To advertise the protection bandwidth for a link that has dedicated 1:1 or shared M:N protection, a link with LPT field Extra Traffic should be advertised. This indicates that bandwidth can be used by LSPs, with the caveat that any LSPs routed over this link will be pre-empted if the resources are needed as a result of a failure over the primary link.

When a failure occurs on a dedicated 1:1 or shared M:N link, the LSPs routed over the link will automatically be switched to the Extra Traffic link that is protecting it.

To support the routing of secondary LSPs for M:N path protection, new extensions must be added to the current GMPLS routing extensions. In particular, there must be a mechanism to advertise secondary bandwidth, and processing rules must be defined for bandwidth accounting when LSP requests arrive at a node. See [17] for a proposal addressing these issues.

Adoption and Implementation of Photonically Switched and GMPLS–Enabled Networks

As with any new technology, service providers will scrutinize photonic switching and control-plane technology before adopting and deploying it in their networks. The evaluation of these new switching systems goes well beyond technological issues. Service providers need compelling business solutions that improve their financial performance. The combination of scalability, transparency, and intelligence offered by GMPLS–enabled photonic switching translates into four key advantages that allow carriers to differentiate themselves from the competition: reduced capital expenditure, operational excellence, revenue growth, and superior network manageability.

Lab and network trials are currently in progress to evaluate both photonic switching and GMPLS control-plane technology. In the first phase of testing, conducted over the last six to 12 months, carriers have primarily concentrated on evaluating the features, characteristics, and performance of MEMS–based photonic switches, typically in standalone configurations. With the critical optical characteristics validated, some carriers are now looking at deploying these switches under control of their existing managements systems. In this model, photonic switches play a functional role similar to that of today's digital cross-connect systems (DCSs), while exploiting the benefits of optics. Once the DCS has efficiently packed the information in the electrical domain and generated SONET–compliant optical signals, the key role of the photonic switch is to dynamically reconfigure the network—at the wavelength level—for restoration or to accommodate changes in bandwidth demand. As mentioned, network reconfiguration is controlled by a centralized network-management system (NMS), which issues switching commands to each switch individually to create optical paths through the networks.

More recent trials are focusing on the GMPLS control plane and more specifically on the solutions offered by GMPLS–enabled photonic switches. The broad interest in GMPLS is illustrated by the fact that incumbent providers and newer entrants alike are planning these trials and even early deployments, demonstrating that the GMPLS control and management plane provides an adequate answer to the yet-unmet need for homogeneous network control. Beyond photonic switches, most trials involve equipment from multiple vendors, encompassing routers, O–E–O switches, metropolitan and long-haul DWDM systems, and many more.

With the choices that need to be made and the difficulty for some vendors to quickly introduce a standards-based control plane, implementations will likely be phased. This allows gradual education of the service-provider community and their customers regarding the advancements and advantages created by the new "Photonic Paradigm." To accelerate creation of best-of-breed multivendor photonic networks, the following step-by-step approach could be considered:

- *Phase I*
  a) Static configuration of an out-of-band control channel (e.g., Ethernet) between the router and the optical switch to demo MPLS interoperability. This supports LSP creation, using RSVP to set up LSPs via the control plane.
  b) Basic LMP procedures between routers and optical switches to demonstrate the functionality of LMP. This phase includes the use of a control channel between the optical switches and bearer channel initialization and verification. It also facilitates the operation of GMPLS by synchronizing the port mappings for the creation of LSPs using GMPLS.

- *Phase II*
  a) This phase concentrates on optical path redundancy and introduces 1+1 network redundancy to provide fault detection and network restoration at SONET time scales. Other restoration mechanisms include 1:N path restoration and M:N bearer channel protection. This phase also introduces load-balanced path selection, an enhancement to the constraint-based path computation to improve load balancing among equal-weight paths.
  b) The final phase adds support for user network interface (UNI) + network-to-network interface (NNI) signaling and introduces LMP–WDM. This interface allows optical switches and routers to communicate with DWDM transport systems, to retrieve health and status information from the optical cloud, and to make it accessible to edge devices for use in network management and path computation.

The evolution to photonically switched networks will be gradual, as several unresolved challenges still lie between today's networks and purely photonic networks. Among the impediments to true optical transparency are all-optical regeneration, optical performance monitoring, and wavelength conversion. In the absence of commercially viable "all-optical" solutions, some measure of optical-to-electronic conversion should be expected in near-term, practical photonic networking architectures. As this is associated with higher cost, the optimal solution will be one that implements the absolute minimum amount of opto-electronics required to ensure the desired optical signal quality and performance monitoring. Concurrently, this solution has to support an easy migration to all-photonic networking once the technological limitations are removed.

The resulting near-term photonic networking architecture is characterized by transparent photonic segments, bounded by opto-electronics. In the medium term, wavelength conversion combined with all-optical regeneration will allow the conversion and regeneration of any bit-rate or data format. Combined with optical signal quality and performance monitoring, this will lead to the implementation of end-to-end photonic transparency.

The opto-electronic boundaries in the near-term architecture will consist primarily of transponders, which allow systems that were not designed to operate in the tightly controlled WDM environment to operate over these systems. Transponders also provide a common interface between multiple vendors' equipment—a necessary provision, as there are many proprietary DWDM implementations. The third function the transponders provide is interfacing to large routers and switches, since OC–48c and OC–192c high-speed interfaces now enable routers and switches to interface directly into DWDM systems.

Currently, interfaces into the WDM systems are SONET–like. The challenge is to determine how SONET–like they will remain. The bit rates will be common to the SONET rates and will have basic SONET framing and multiplexing. Whatever SONET functions remain, all parameters need to be configurable to allow service providers increased flexibility in their migration strategy. Some service providers will continue to build their SONET network in the same way but will use a photonic switch to interconnect rings. This is a very easy step to take as it preserves the existing ring architecture and restoration techniques (diversely routed 1+1 capacity and SONET ring protocols), while it offers immediate benefits such as simplified service provisioning and management. Additionally, the combined cost of photonic switches and transponder-based DWDM is much lower than that of ADMs and transponder-based DWDM.

An intermediate step toward photonic networking is to build SONET rings over a mesh infrastructure. In the final stage, SONET can gradually migrate to the edges of the network with the photonic switches in the middle, providing the routing and restoration functions. The aforementioned implementations can coexist with existing SONET rings while offering benefits like "point-and-click" provisioning and the ability to introduce class-of-service (COS) networking, allowing a service provider to offer traditional SONET services with 50 ms restoration, mesh reroute restoration services, and unprotected circuits, all on the same network. As these services would be provided at various price levels, customers can balance service quality and cost.

The opto-electronic-bounded photonic network combines the benefits of photonic switching with advances in DWDM technology to create a short-term network architecture that delivers managed multigigabit bandwidth and provides reliable, wavelength-level, traffic-engineered network interfaces to the service platforms. The service platforms include new high-capacity IP routers, O–E–O switches, ATM switches, and SONET ADMs, which are redeployed from the transport to the service layer. The service layer delivers services to users in various forms and relies entirely on the photonic layer for the delivery of multigigabit bandwidth where and when it is needed. SONET transport gives way to photonic transport, and bandwidth is provisioned, not at TDM granularities, but rather at wavelength granularity. To meet exponential growth rates, rapid provisioning is an integral part of the new architecture. While initial implementations of this model will support error detection, fault isolation, and restoration via SONET, restoration will gradually move to the photonic layer.

The premise of photonic networks requires the availability of tools to measure and control the smallest granular component of such networks—the wavelength channel. These functions include the monitoring of amplifiers and switches at add/drop sites, the deployment and commissioning of DWDM routes, and the restoration and protection of networks. This must be accomplished with speed and accuracy and implies that performance measurements have to be done while keeping the channel in optical format. As photonic networks mature, it will be necessary to generate a more detailed picture of the channel "health" in a manner that can be communicated to the network control entities as well as between other network elements. Detailed performance monitoring at the "digital level" will be performed at the service interfaces.

Conclusions

While today's optical networks are highly reliable and carry remarkable amounts of bandwidth, the optical layer of these networks is fundamentally static and provides for only simple point-to-point transport. All channel management, restoration, protection, provisioning, routing, and switching of optical channels is done in the electronic layer. This requires hundreds of millions of dollars in inflexible O–E–O conversions to accomplish basic bandwidth-management functions. New photonic switches, ultra-long-haul DWDM equipment, optical add/drop multiplexers (OADMs), and photonic-layer protection and restoration functions are becoming available to move functionality from the electronic to the optical layer, thereby dramatically reducing the cost of optical networks.

GMPLS will constitute an integral part of next-generation data and optical networks. It provides the necessary linkage between the IP and photonic layers, allowing interoperable, scalable, parallel, and cohesive evolution of networks in the IP and photonic dimensions. The functionality delivered by GMPLS allows network operators to scale their networks well beyond current limitations implicitly created by the segregation of the transport network from the data above. The signaling capabilities of GMPLS will allow service providers to quickly build out high-capacity agile infrastructures that support fast provisioning of connection services. Now that demand for data transport has eclipsed that for voice, it is appropriate that the control-plane mechanism expand to embrace data transport more closely while still serving the incumbent needs of voice transport. Service providers can incrementally deploy GMPLS–based products in existing networks to decrease costs without impacting service quality. Furthermore, the flexible M:N protection and restoration capabilities of GMPLS allow efficient addressing of network survivability, while opening the door to new types of services.

From a service provider's perspective, GMPLS enables three levels of advanced management and control. At the lowest level, flexible and rapid (optical) bandwidth creation and distribution is enabled through signaling, routing, and the link management of wavelengths. At the second level, service creation and management is enabled through the same control plane but now using timeslots, packets, and cells. Since both layers utilize a single set of semantics, service creation, deployment, and provisioning can be accelerated, while the common network knowledge shared between service and transport equipment facilitates a more efficient and faster service restoration, load balancing, and path-performance optimization. At the third level, the GMPLS framework can potentially provide the first unifying management plane for data networks, similar to what signaling system 7 (SS7) does for voice networks. GMPLS allows homogeneous network control, which is not possible today. Substantial impact will be derived from having network-management applications leverage the control plane to collect and correlate service-related data, which can then be applied to applications such as service scheduling, service-level agreement (SLA) monitoring, network-utilization analysis, traffic-flow patterns, connection trending analysis, and planning tools. Through GMPLS–enabled access to network information, operations support system (OSS) and other software vendors can build vendor- and platform-independent applications that run on standalone machines, while the service provider only has to give these applications access to the link-state database (LSD) to bring them on-line. Carriers will greatly benefit from a common set of management semantics to unify heterogeneous optical networks and deliver consistent information across all elements. Lack of such unified management information throttles today's optical networks, limiting performance, cost-effectiveness, and the carriers' freedom to choose best-in-class elements.

Given the challenge of quickly introducing a standards-based control plane, implementations will likely be phased. Calient was preparing to co-launch GMPLS Test Bed activity in 4Q 2001, involving a range of DWDM, gigabit router, and photonic switching players. The goal is to prove that wavelength routing and connection management can be achieved across multiple network elements from multiple vendors. This effort was to be done in concert with the Internet Engineering Task Force (IETF) committee work and constitutes one of the first cross-industry and cross-element collaborative standards implementation efforts.

References

[1] Rajiv Ramaswami and K.N. Sivarajan, *Optical Networks – A Practical Perspective* (Morgan and Kaufmann, 2nd Edition, 2001), Chap. 7.

[2] Mohan-Rao G. Lingampalli and Mohcene Mezhoudi, "Temperature Sensitivity of Lithium Niobate Optical Switches," Proceedings of SPIE Conference, Vol. 2290 (San Diego, Calif.: July 1994), pp. 133-138

[3] D. Blumenthal, "The S3 of MEMS–Based Photonic Switches: Scalability, Size, and Serviceability," *Fiber Optics Online* (November 2000).

[4] Daniel O. Awduche et al., "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Cross Connects," Internet Draft, Work in Progress (April 2001), draft-awduche-mpls-te-optical-03.txt.

[5] J. P. Lang et al., "Link Management Protocol (LMP)," Internet Draft, Work in Progress (September 2001), draft-ietf-ccamp-lmp-01.txt.

[6] K. Kompella et al., "LSP Hierarchy with MPLS TE," Internet draft, Work in Progress (March 2001), draft-ietf-mpls-lsp-hierarchy-02.txt.

[7] A. Banerjee et al., "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements," *IEEE Communications*, Vol. 39, No. 1 (January 2001).

[8] Chiu, A., Strand, J., and Tkach, R., "Unique Features and Requirements for the Optical Layer Control Plane", Internet Draft, Work in Progress, draft-chiu—strand-unique-OLCP-01.txt.

[9] Fredette, A. et al., "Link Management Protocol (LMP) for WDM Transmission Systems," Internet Draft, Work in Progress, draft-fredette-lmp-wdm-00.txt.

[10] Lang, J. P. et al., "Link Management Protocol (LMP)," Internet Draft, Work in Progress, draft-ietf-mpls-lmp-02.txt.

[11] Ashwood-Smith, P. et al., "Generalized MPLS Signaling – RSVP–TE Extensions," Internet Draft, Work in Progress, draft-ietf-mpls-generalized-rsvp-te-03.txt.

[12] Ashwood-Smith, P. et al., "Generalized MPLS – Signaling Functional Description," Internet Draft, Work in Progress, draft-ietf-mpls-generalized-mpls-signaling-04.txt.

[13] Kompella, K. and Rekhter, Y., "Multi-Area MPLS Traffic Engineering," Internet Draft, Work in Progress, draft-kompella-mpls-multiarea-te-00.txt.

[14] Awduche, D. et al., "RSVP–TE: Extensions to RSVP for LSP Tunnels," Internet Draft, Work in Progress, draft-ietf-mpls-rsvp-lsp-tunnel-08.txt.

[15] Kompella, K. et al., "OSPF Extensions in Support of MPL(ambda)S," Internet Draft, Work in Progress, draft-kompella-ospf-gmpls-extensions-01.txt.

[16] Kompella, K. et al., "IS–IS Extensions in Support of Generalized MPLS," Internet Draft, Work in Progress, draft-ietf-isis-gmpls-extensions-0.2txt.

[17] Kompella, K., Lang, J.P., and Drake, J., "Bandwidth Accounting in Support of Secondary LSPs," Internet Draft, Work in Progress.

Architectures for Metro Optical Networks

Joon Choi
Director, Product Management
Mayan Networks

Although the prevalent metropolitan optical networks (MON) consist of the traditional synchronous optical network (SONET) architecture, changes in the MON are imminent. This paper discusses the evolution and challenges of competitive local-exchange carriers (CLEC) and service providers (SP), evaluates existing solutions, and examines the requirements for addressing the emerging challenges.

Evolution

The Telecommunications Act of 1996 was developed to foster more competition in the metro area and in local-loop access. The Act also introduced the first set of metrics used to evaluate the success of CLECs and SPs. The network elements (NE) addressed by the first metrics were the number of fiber miles to which a CLEC or SP had access and the number of buildings to which the fiber had access.

SPs began to provide service based on those NEs, including data services and digital subscriber line (DSL), that were very popular with providers. Eventually the CLECs and SPs formed partnerships with application service providers (ASP) in offering Internet protocol (IP) services. Next they focused on increasing capabilities and seeking ways to bill for service on a usage basis in lieu of offering a monthly flat-rate service charge.

Now the following questions must be addressed:

- How do CLECs and SPs evolve further?
- What are their challenges in terms of the metro edge market?
- How will those challenges be addressed?

Challenges

Some of the solutions addressed by a study in early 2000 were the optical IP, multiservice dense wavelength division multiplexing (DWDM), and the multiservice SONET. The numbers were quite high in terms of projections for 2001 to 2004. On the multiservice DWDM, for example, the number doubled from 2001 to 2002. The multiservice SONET showed a market increase to approximately $6.3 billion in 2004.

Admittedly, this research was done when the stock market was in an exuberant state, so the numbers may have been on the high side, but the following conclusions can be drawn nonetheless:

- Various solutions exist in the metropolitan-area network (MAN).
- SONET as an infrastructure will continue to exist.
- SONET must evolve to provide the performance required by, and meet the criteria of, multiservice.

Current Edge Architecture

At the metro edge, where access exists to the SONET access ring in the MAN, there are multiple service types, such as the following:

- A voice network that requires a fairly delayed, sensitive type of traffic
- A service offering virtual private networks (VPN)
- An enterprise system connection (ESCON)
- Web hosting

Physical interfaces, which are required for the access, are also found at the edge:

- Ethernet interfaces
- Telephone company interfaces; e.g., T1 and T3
- DSL or direct optical interfaces; e.g., optical carrier (OC–3, OC–12, and OC–48)

Riding on the physical interfaces are the protocols for handling various functions:

- A pass-through of time division multiplex (TDM) traffic
- A pass-through of asynchronous transfer mode (ATM)
- IP
- Frame relay (FR)

SONET Optimization

SONET was originally designed to be used for voice, so it is optimized for voice and based on TDM. Consider the case of an SP that offers three megabit (Mb) Ethernet service and provides direct access to the SONET access ring, in this case an OC–3 metro ring. The SP utilizes a legacy add/drop-mix type of equipment which will pass the traffic through, mapping the three Mb service directly onto the granularity of the SONET payload—or the time slot which, in this case, is a synchronous transfer signal (STS)-1 running at 51 Mb per second (Mbps).

Only one third of the entire SONET ring is allocated for three Mb of data stream, the equivalent of a six percent utilization—not an optimum usage of SONET carrying Ethernet directly. Additional grooming and, perhaps, some additional intelligence are required to best optimize the SONET to handle the ever-increasing data-oriented traffic

Solutions

Various solutions, including SONET equipment, ATM virtual path rings (VPR), passive optical networks (PON), wavelength division multiplexing (WDM), dense WDM (DWDM), and gigabit Ethernet (GbE), are currently available. Each solution has its own advantages and issues that must be addressed.

SONET Equipment

Traditional SONET equipment consist of muxes and digital cross-connect (DXC) that all operate at Layer 1. Because there is no higher-layer intelligence, there is no optimization for the data-oriented traffic via the existing SONET infrastructure. SONET also does not scale well and, in the past, the equipment had to be removed and replaced with other types of equipment that might provide better service.

ATM Virtual Path Rings

The ATM VPRs represent another solution, wherein a service layer is added onto the existing SONET infrastructure. Adding a protocol such as ATM provides some guaranteed service, a good quality of service (QoS), and a mixture of different types of traffic, but it also increases overhead.

The biggest issue with the VPR solution is that it is of proprietary nature at this time. Each vendor is required to test and address interoperability with other vendor's equipment.

PON

PON can be a good, cost-effective solution, and it has become popular in the campus-area network as well as the enterprise. Today, however, the PON is still a proprietary solution that requires end-to-end products from a single vendor. The PON also does not offer any edge intelligence; there is no higher layer of processing, such as IP or ATM.

WDM

WDM edge switches and the DWDM edge do aggregate traffic onto the backbone and expand the capacity of the fiber. However, in terms of cost, as the DWDMs in the long-haul area move closer toward the edge and within the access of the MAN, the current price of equipment—not including the bandwidth (BW), at least in nominal cases—is nevertheless still too high for the metro-area edge in comparison with alternative solutions.

Both WDM and DWDM are composed of proprietary solutions, requiring products from either a single vendor or, with a limited interoperability, multiple vendors.

GbE

The GbE started from the local-area network (LAN) side and provides a very good price for BW connectivity. But extending the GbE to the metro access and edge is practical only if the price can remain the same and the product can continue to offer carrier-class reliability and resiliency, which remains a challenge.

The Smarter Metro Edge

The smarter metro edge has four basic ingredients: interoperability, grooming, transport, and provisioning and offer solution to the challenges found in the MON

Interoperability

As the SONET continues to grow and evolve, it is critical for new equipment to be interoperable with the existing equipment, not only at the physical layer but at the signaling and control layers as well. Some equipment must be managed and operated through individual vendors, and multivendor equipment must be managed and operated through the data communications channel (DCC) and other standards-based channels.

Smarter Grooming

One aspect of smarter grooming is providing the DXC all the way down to digital signal (DS)–0 to provide better visibility into the lower-level traffic, which might range up to OC–48.

The other aspect is the ability to provide a distributed grooming capability where, at the entrance of the SONET access ring, the DXC capability can be provided, thereby packing the SONET ring BW more tightly and providing a cost savings.

SONET Transport

To carry data over SONET, the BW must be maximized and efficiently utilized. A small amount of intelligence should be added, while maintaining interoperability and standards over the SONET ring and continuing to provide the maximum utilization of that data BW riding over SONET.

Smart Provisioning

Decoupling the physical interfaces from the protocol or service aspects should enable the provisioning of services on a per-port basis. With a T1 module, for example, each port should be able to provide TDM, ATM, FR, and so on, something only possible with decoupling. A distributed architecture is required for smart provisioning.

DSL and Remote Terminals

Jon Cordova

Directing Analyst, WAN, Internet, and New Metro Access

Infonetics Research

This article discusses digital subscriber line (DSL) and remote terminals, market drivers, deployment models, and market forecasts.

Background on Remote Terminals

Remote Terminal Basics

Remote terminals (RT) are small, weatherproof vaults, huts, or cabinets that were originally used to house digital loop carrier (DLC) systems to provide voice services and pair-gain and to avoid copper home runs. They are usually located within a mile of the final drop to a residence or a small business. Why should RTs be used for DSL deployment? Primarily because the central office (CO) is, in many cases, too far from the user; DSL typically has a three-mile reach from the CO to the customer. In addition, the regional Bell operating companies (RBOC) can leverage their existing DLC RT infrastructure. RTs currently house legacy and generic requirement (GR) 303 next-generation digital loop carriers (NGDLC) and will house new-generation broadband loop carriers in the future. The new-generation broadband loop carriers integrate various combinations of plain old telephone service (POTS), DSL, media gateway, and loop-management functionality.

Remote Terminal Examples

Some examples of RTs are shown in *Figure 1*. The first picture at the upper left is a remote hut. Below that, is a photo of a small vault. The digital subscriber line access multiplexer (DSLAM) at the top right is adjacent to a DLC. Finally, at the lower right is an ordinary DLC with a module for upgrading to DSL services.

Industry Trends

The fundamental consumer demand for DSL remains strong. In North America, DSL providers grew 20 to 40 percent during the first quarter of 2001, ending with 3.9 million subscribers. Cable subscriptions grew 15 to 25 percent, ending with 5.6 million subscribers for the same quarter. DSL providers maintain subscriber backlogs, and prices are increasing in accordance with demand.

Still, a large percentage of households cannot get DSL. The DSL coverage rate for BellSouth is about 60 percent of its customer base, and for SBC it runs at about 50 percent. Verizon and Quest are about the same, at 40 percent. Currently, RTs for the RBOCs are estimated at over 100,000 units. These RTs are serving about one third of residential DSL lines and will probably serve over 50 percent in the future.

Market Drivers

Residential Drivers

For residences, some of the market drivers are high-speed Internet access, video-on-demand, voice over digital subscriber line (VoDSL), distance learning, and home networking.

Business Drivers

The business drivers are telecommuting with virtual private networks (VPN), videoconferencing, VoDSL, distance learning, and, finally, asymmetric digital subscriber line (ADSL) adoption by small offices/home offices (SOHO) and small businesses for cases where ADSL speeds are sufficient to meet their business requirements.

Deployment Models

CO Deployment Model

The most prevalent deployment model in the CO has a DSLAM and a DLC both located at the incumbent local-exchange carrier (ILEC) facility, and the subscriber receives services through a standard DSL modem or router plus a phone set, or, in some cases, a residential gateway (see *Figure 2*). The signal is split at the CO, with the data traffic being sent to the DSLAM and the voice traffic to a local switch.

Next-Generation DLC Deployment Model

Another deployment model is that of the NGDLC located in an RT (see *Figure 3*). On the left side, the subscriber picture looks about the same. On the right side, the data and voice traffic is split and then sent to the respective data and voice networks. NGDLCs maintain separate ports for POTS and DSL with the splitter typically integrated on the line cards.

Remote DSLAM Deployment Model

The remote DSLAM deployment model calls for both the DSLAM and the DLC to be located in a RT (see *Figure 4*). Here, there is a full-service DSLAM working in conjunction with a DLC, so the ADSL and POTS lines need to be split at the cross-connect before voice and data traffic can traverse their respective public switched telephone network (PSTN) and asynchronous transfer mode (ATM) networks.

FIGURE 1

Examples of RTs

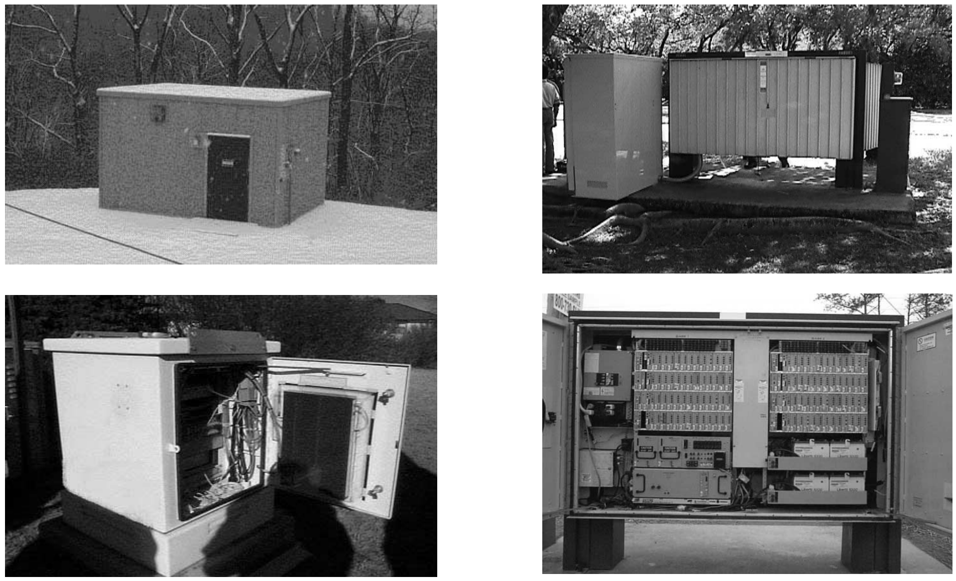

Future Converged Packet-Based Public Network

Of course, the networks' utopia is the future converged packet-based public network (see *Figure 5*). In this case, one broadband loop carrier in the CO or the RT aggregates and distributes all the traffic. On the left side is the typical subscriber customer-premises equipment (CPE). The unique characteristics of these broadband loop carriers are their integrated port and line-card functionality, which includes DSLAM, DLC, media gateway, and loop-management functionality. They require no internal or external splitter, help eliminate truck-rolls, and can be remotely provisioned. Many NGDLCs vendors will likely start to build all of the aforementioned functionality into their products, and both CO and RT deployments will be popular. Broadband loop

FIGURE 2

CO Deployment Model

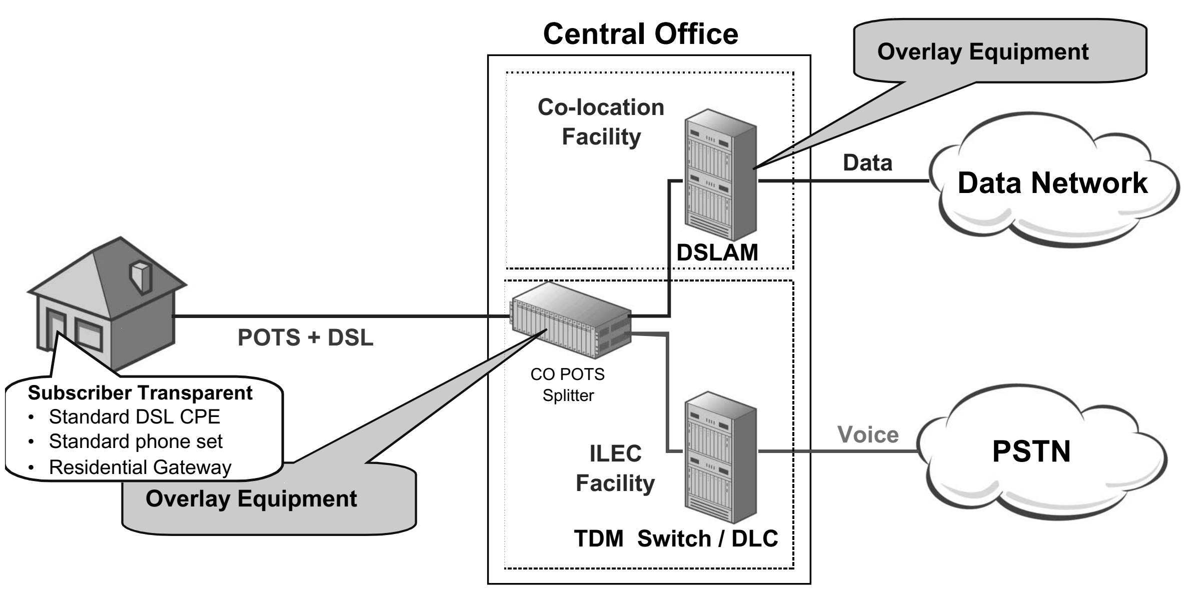

FIGURE 3

NGDLC Deployment Model

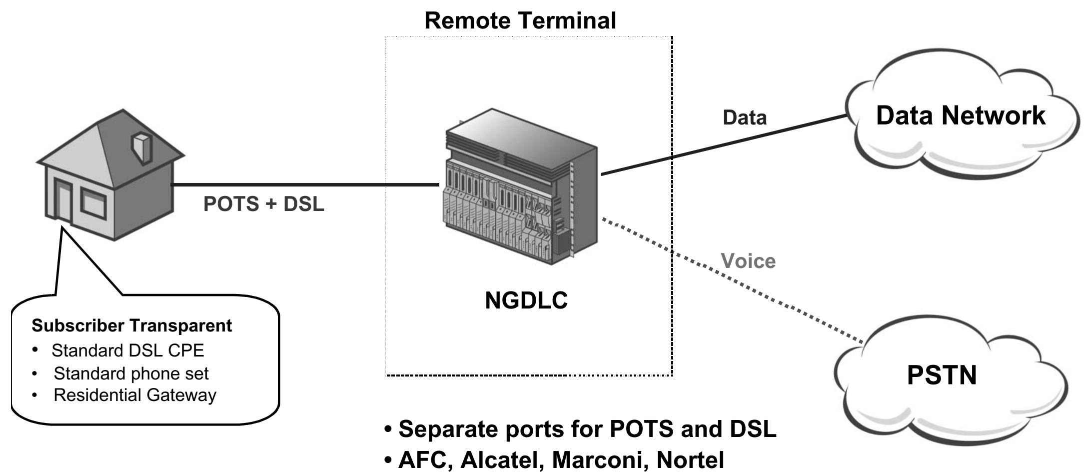

carriers have not gained much traction yet, but are currently in trials with a number of the RBOCs.

Market Forecast and Predictions

Worldwide DLC DSL Line-Card Revenues
Finally, *Figure 6* deals with the market forecast and some predictions for worldwide DSL line-card revenues. This market is expected to grow from about 300 million at the end of 2001 to about 750 million by the end of 2004. However, this does not include the broadband loop-carrier revenue, nor does it include any chassis revenue or revenue from deployments in COs. It is a very simple, realistic forecast.

Forecast Factors and Assumptions
Some of the factors and assumptions used in the forecast are important to understand.

- SBC has approved Alcatel and AFC platforms for Project Pronto, which aims to cover about 80 percent of the customer base with over 17,000 RTs by 2003.

FIGURE 4

Remote DSLAM Deployment Model

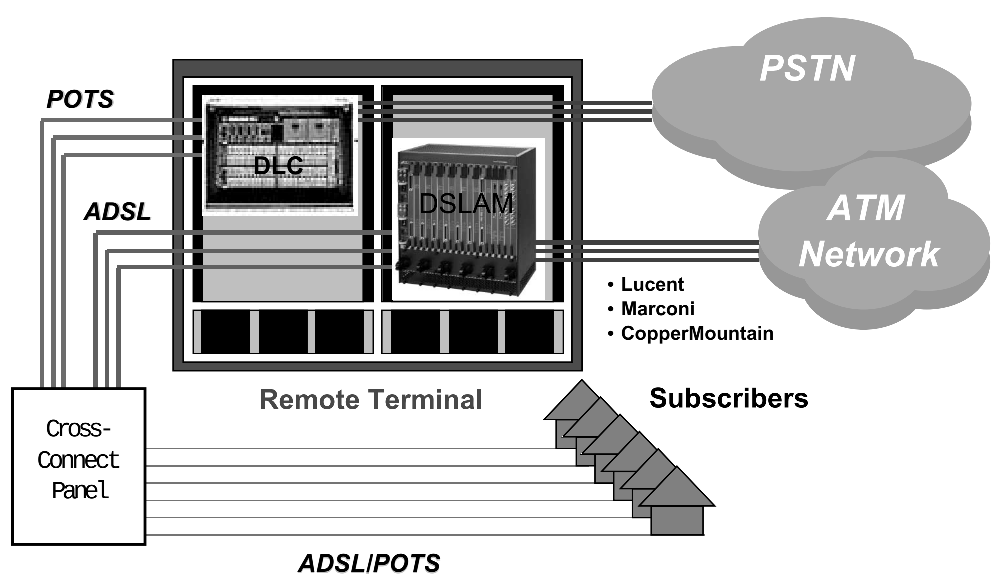

FIGURE 5

Future Converged Packet-Based Public Network

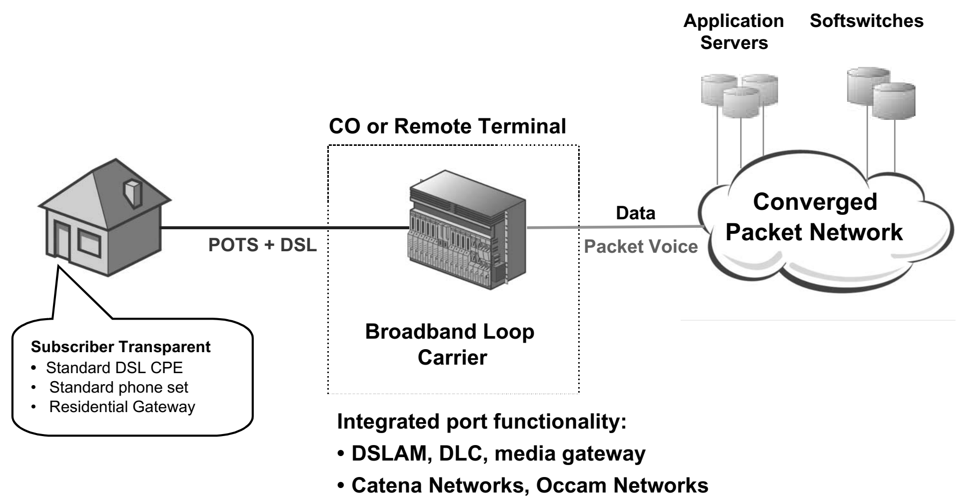

- Lucent announced an agreement to provide Stinger DSLAMs in Qwest RTs that will add about 6 million DSL–capable customers by the end of 2002.
- BellSouth and Verizon are a little bit more on the fence. They have talked about some different deployment models, and BellSouth has indicated it will deploy Marconi DSL remotes in 2001 that would add 100,000 new DSL–ready households per year. Verizon has reported interest in AFC, Alcatel, and Lucent solutions, but has not committed to remote DSL deployments because of unbundling requirements and uncertainty about how regulatory legislation will unfold.
- The last, and most important, factor is that the RBOCs will have much more incentive to modernize their access networks and gain return on investment if the Internet Freedom and Broadband Deployment Act of 2001, which is also termed the Tauzin-Dingle Bill, is passed. If it is not, there will probably continue to be deployments by the RBOCs, with waivers in some regions and some competition from the competitive local exchange carriers (CLEC) and the inter-exchange carriers (IXCs) as well. However, this bill is key to the RBOCs getting very aggressive with their deployment plans.

FIGURE 6

Worldwide DLC DSL Line-Card Revenues

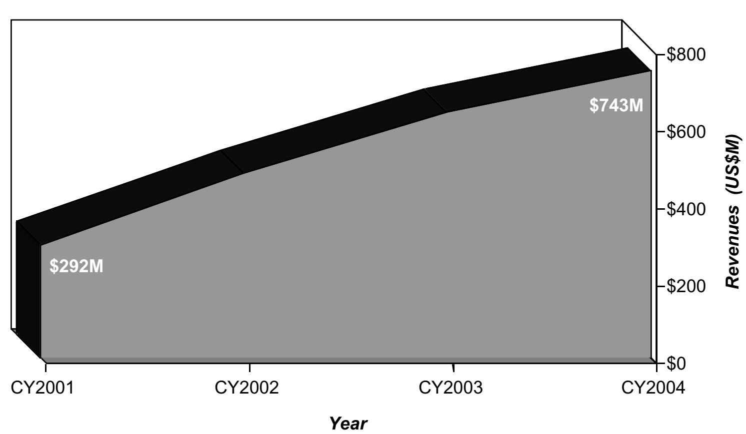

Source: *Broadband Hardware: IADs, DSL, and Cable Quarterly Worldwide Market Share and Forecasts*

After the DSLAM: From CO Magic to Service Provisioning

Ed Cox
Director, Business Development
Santera Systems

This article discusses changes inside the central office (CO) and what needs to happen within the switching environs for broadband services to be delivered while supporting the legacy voice work.

Disruption in the Service Provider's World

A lot of things are going on in the market that cause disruption in the service provider's world, particularly within the last year (see *Figure 1*). Technology is changing rapidly. Service providers have the ability to introduce new things and implement new services, but they are saddled with old architectures, meaning that the previous generation of technology has a shorter life cycle, putting more strain on the network's ability to absorb it. Standards and protocols are constantly changing, making more customer and employee training necessary. The market is disrupted by missing market-entry barriers, bandwidth hunger, and challenges caused by convergence. There is widespread confusion as well as competition. And of course, in all of this there are the financial issues that go with the more rapid adoption of technology—how to use it, how to quickly recover from it, and what services can be provided while it is in the network.

Access Demand Cannot Find the Network

The access network continues to grow and diversify (see *Figure 2*). There is bandwidth improvement within the core of the transmission network as well, whether it is digital subscriber line (DSL) or cable modems on the access side, or any of the dense wavelength division multiplexing (DWDM) types of things that are being done in the core. But nothing has happened to date in between those two, which is the switching equipment that connects all of that high-bandwidth access with all the high-band transmission work going on.

Public-Switching Evolution

The public-switching evolution started out in the 1900s, and about every 25 years or so there have been major changes (see *Figure 3*). The last major change that happened within the switching environment was the move from analog to digital. That has lasted the telecom industry for about 30 years, but in the move from the circuit-based world, with pure voice, to the packet-based world, where voice is just one of many different application bases, some major changes need to happen. IBM learned this lesson very painfully in the move from a monolithic computing model to a distributed model. The question, then, is whether, as they move from a circuit- to a packet-based world, the legacy vendors and carriers are going to realize how upside-down current business practices are actually going to leave them.

Size of United States Replacement Market

There are a number of switches from a voice perspective, with over $50 billion in replacement value, that are either discontinued with no support or in their last year of support (see *Figure 4*). If one looks at the numbers and starts to add them up, even with the next-generation switching dollars, there is a lot of change that needs to happen within the next 5, 10, and 15 years. There is a large amount of switching equipment out there that has been the workhorse of the service-delivery model for the voice-only world. Now, with the move into a packet world, not only is different equipment needed, but the equipment that the industry has relied on is not quite as confident a continuation.

Transition Network

The current transition network (see *Figure 5*) is introducing a variety of different gateways to the softswitch to control those gateways that handle the operation and interoperation between the different routers, asynchronous transfer mode (ATM), frame switches, and the digital subscriber line access multiplexers (DSLAM) themselves; the different voice types of equipment (seen on the right hand side between the Class 5 and the Class 4 switch); digital loop carriers (DLC); and the like. This is, of course, introducing many boxes with many networks and many headaches. The more network elements involved in providing service, the more networks a carrier has to maintain and the more difficult it is for them to move quickly, nimbly, or maybe even at all and the more difficult it is for them to offer those services in some kind of a cost-effective manner.

FIGURE 1

Disruption in the Service Provider's World

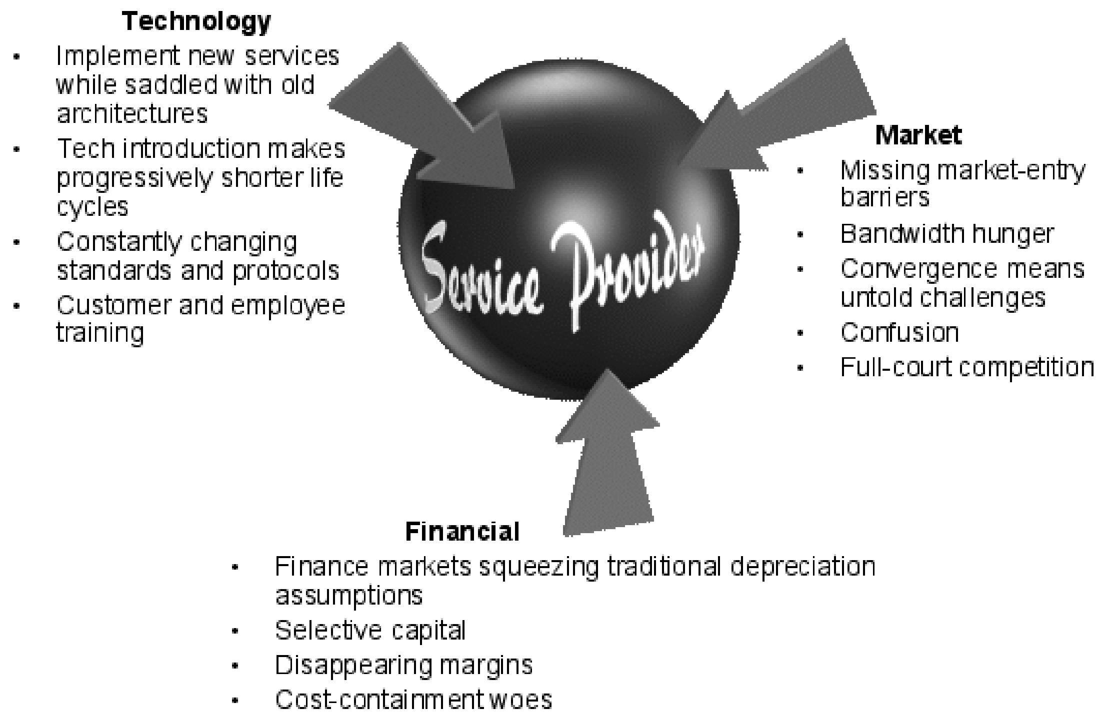

FIGURE 2

Access Demand Can't Find the Network

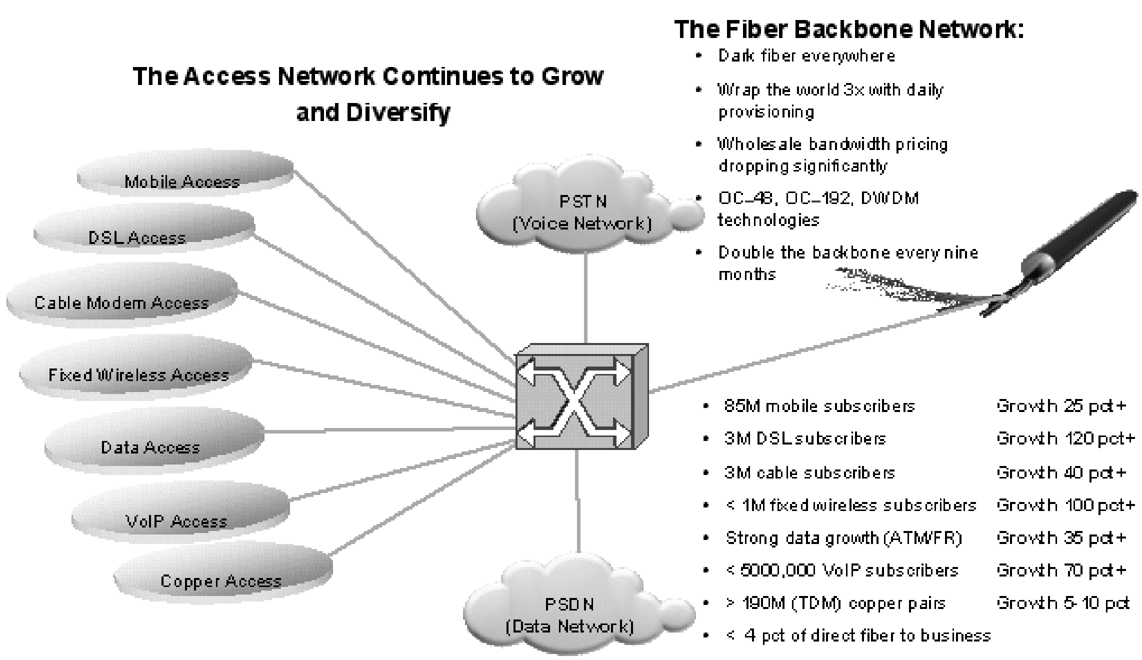

FIGURE 3

Public-Switching Evolution

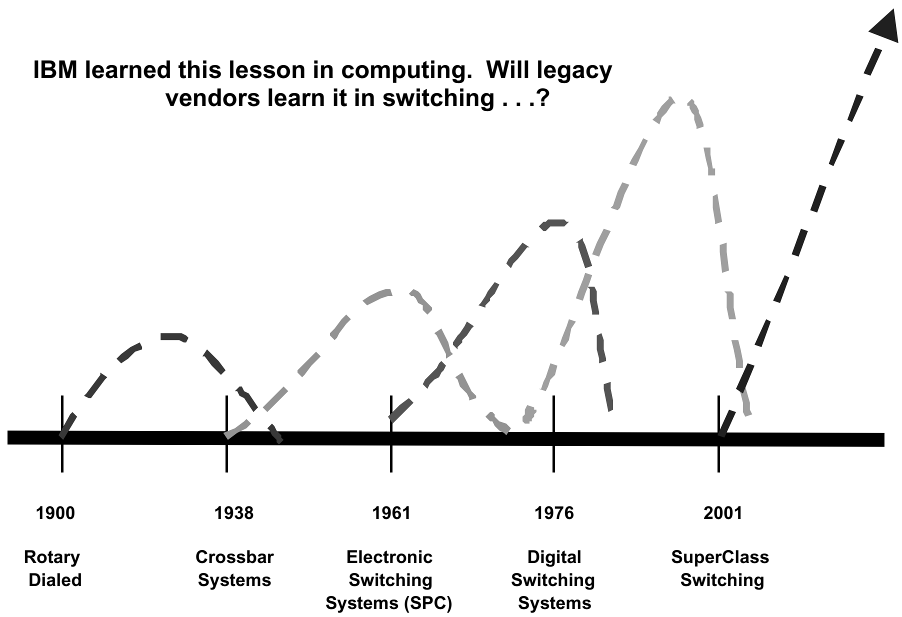

Integrated Network

The fully integrated network (see *Figure 6*), when it arrives, will have the ability to put the softswitch and the media gateway together as a superclass switch, or the next switching element, so that all the pieces and information flows that come from the optical-transport network can be interconnected and so that they get down to the customer-premises equipment (CPE) with the switching, adaptation, and services treatment that needs to happen. In doing this combination of functionality, many features will have multiple functions and many opportunities as the service-delivery equation is simplified greatly.

Sustainable Architecture

This needs to be done on a sustainable architecture, where the media gateway is able to provide a variety of traffic types in and out, with the ability to scale and work within different types of bandwidth requirements with different applications in the broadband network. That media gateway will be controlled by the softswitch, which has continued to increase in its definition and importance within the network. At one time the media gateway only included the call control. Now it not only controls the media gateway controller (or softswitch or call agent) but also serves as the repository to run the software functions of the signaling gateway; the application server, if there is one on board; and the element-management system in order to work with all of the different provisioning and management systems. Third, because no single vendor is going to be able to write all the software in the telecom space, the platform needs to be able to provide access to third-party application servers and to continue to increase the value and the number of services that are provided and supported by that particular platform (see *Figure 7*).

DSL Services: Current Transition Network

DSL services in the transition network are a complex challenge (see *Figure 8*). There are a number of network elements involved in providing the service, whether extracting the voice off of a voice gateway or using a plain old telephone

FIGURE 4

Size of United States Replacement Market

Switch	Status	Number of NA Switches
Ericsson AXE–10	Discontinued No Support	Over 80 switches
Lucent 1A	Discontinued No Support	>300 switches
GTD–5	Last Year of Support	>1,900 switches
Siemens EWSD	Current Engineering Only	>1,600 switches
Nortel DMS 100	Last Year of Support	>8,300 switches
Lucent 4ESS	Current Engineering Only	140+ switches
Nortel DMS250	Current Engineering Only	>950 switches
Lucent 5ESS	Current Engineering Only	10,000+ switches
Nortel DMS–10	Discontinued No Support	>4,800 switches

Over $50B in Replacement Value

FIGURE 5

Transition Network

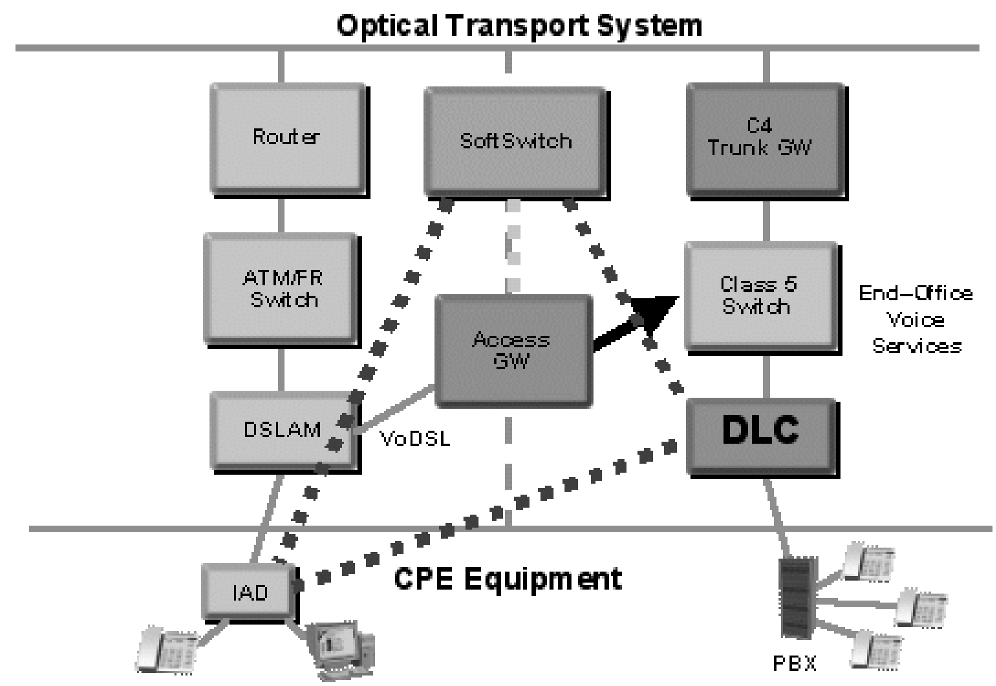

FIGURE 6

Integrated Network

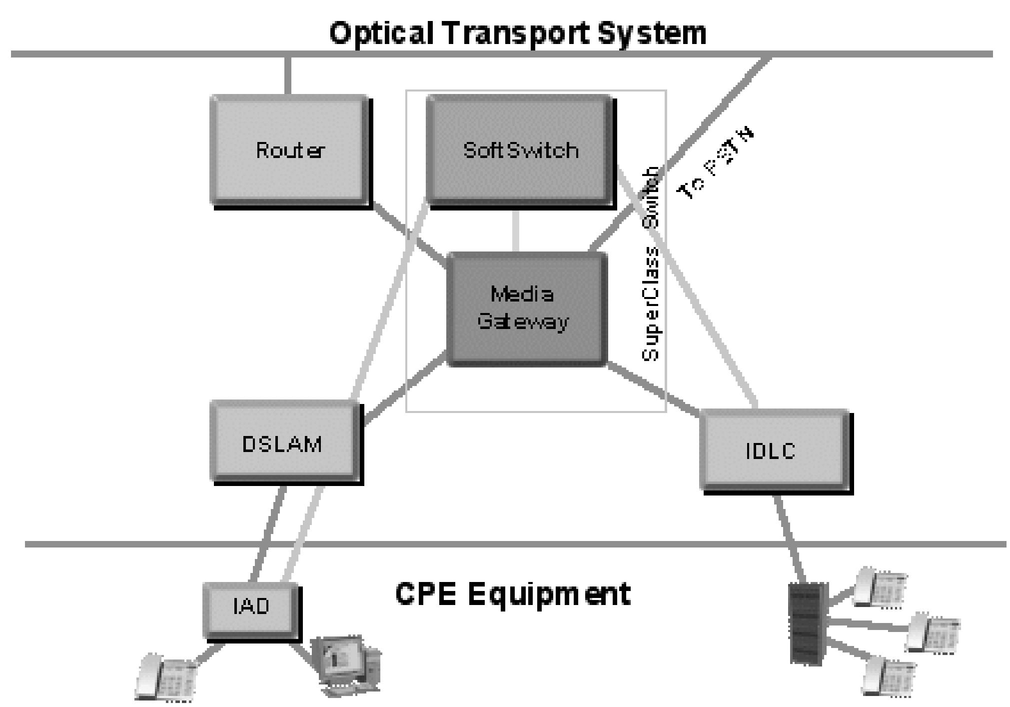

FIGURE 7

Sustainable Architecture

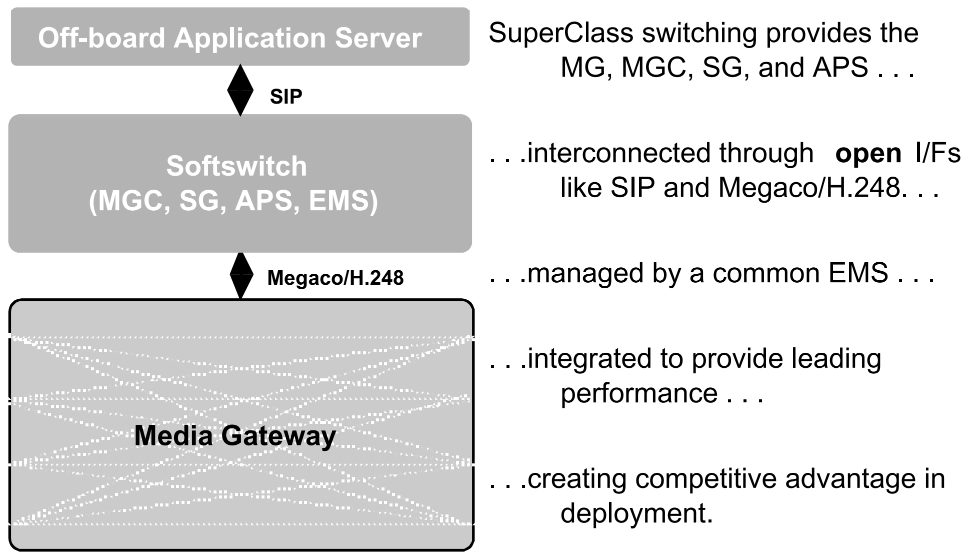

service (POTS) splitter to communicate with the DSLAM and/or the data switch. A long-term service option to supply 7 million lines of voice-over-broadband servers in a single metropolitan area really is a difficult situation to manage. It is a complexity as far as management is concerned. It has lower reliability because it has so many points involved in the total end service and a higher cost to provision.

DSL Services: Integrated Network

Ultimately, next-generation network switching needs to provide an aggregation of functions and to take the traffic directly in from the DSLAM. It needs to take that traffic and switch it, provide services treatment at wire speed, and, depending on the type of service and how that line is provi-

FIGURE 8

DSL Services—Current Transition Network

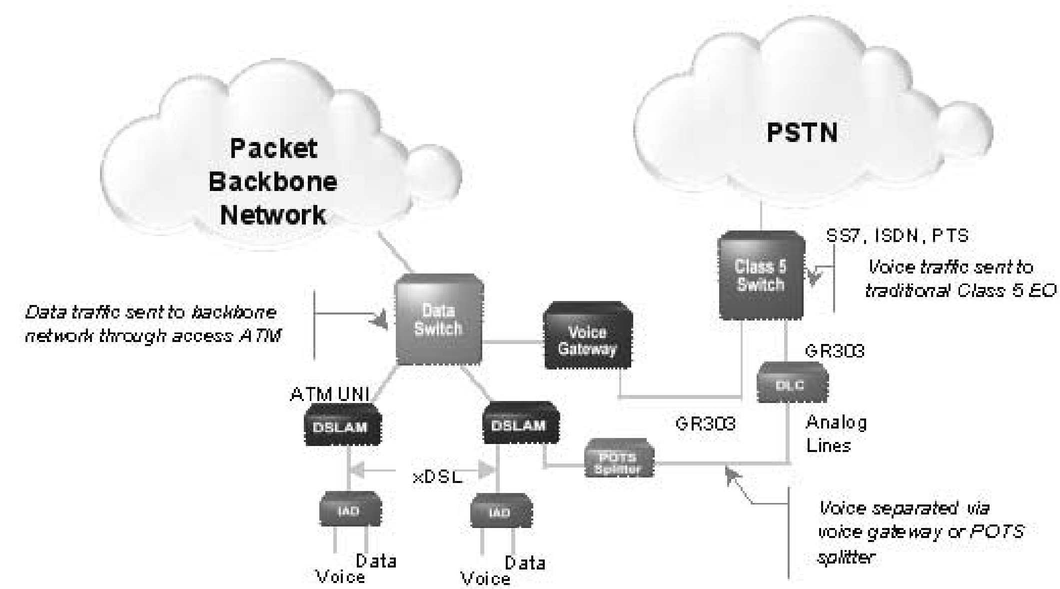

sioned, to determine whether that traffic then needs to be switched off onto the packet backbone network (regardless of whether it is frame relay, IP, or ATM based) or whether to go back into the legacy environment and continue to support the backbone or the POTS world with equal ability (see *Figure 9*).

Putting Brawn and Brains in the CO

Brawn and brains need to be put back into the CO. DSL is just one example of what next-generation switching must do in order to serve as the switching vehicle over the long term, like the time division multiplex (TDM) switching has done for the last 40 years. This next-generation platform must be put in the network with the knowledge and confidence that it can be counted on for a great period of time. As the industry moves to the packet world, the legacy switching is either dying or dead. The number of switches that are in their last year or have been discontinued in terms of service and support may leave some vendors aching for a cure—but nevertheless, CO switch expectations need to increase. The vendors providing this space need to deliver to that new level of expectation. The industry needs to set the hurdles for vendors to be able to jump over.

The decision then becomes, if vendors are actually performing to spec, should the industry look at a point solution that solves the problem for today by switching within transition networks to solve bottlenecks? Or should it look at a platform decision to really get at the switching needs from a long-term perspective, supporting graceful migration by switching within integrated networks? If it does that, the industry will be able to provide a platform that allows the carrier a variety of different vehicles.

Greenfield, replacement, or cap and grow all follow same architectural design—SuperClass switching. This new SuperClass-switch concept can support the voice as well as the data world, and all the services and points in between.

FIGURE 9

DSL Services—Integrated Network

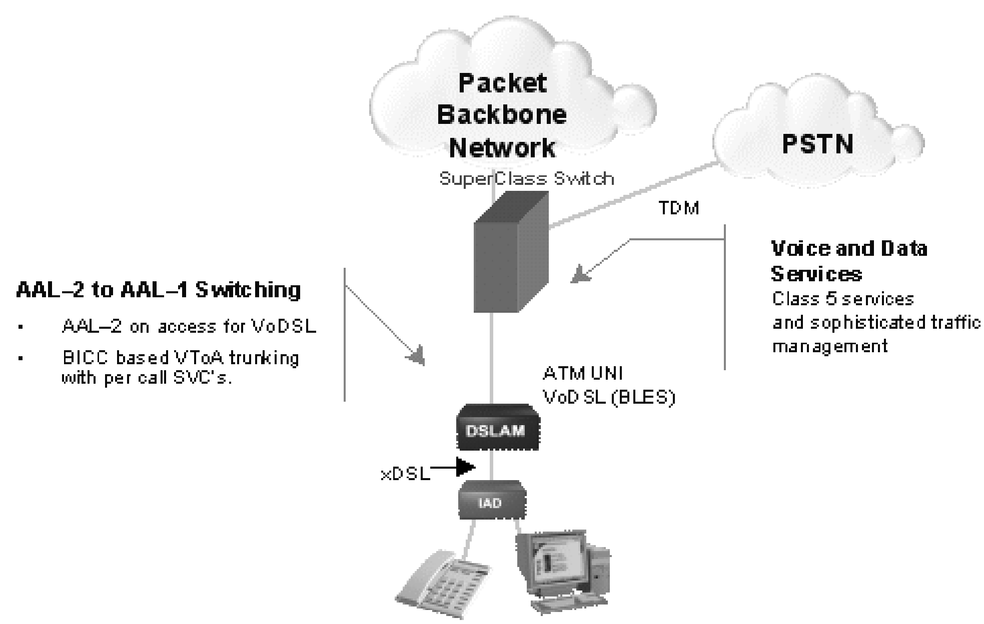

Physical-Layer Considerations for Advanced TDMA CATV Return Path

Bruce J. Currivan
Director of Systems Architecture, Broadband Systems Engineering
Broadcom Corporation, USA

Thomas J. Kolze
Senior Principal Scientist
Broadcom Corporation, USA

Jonathan S. Min
Principal Scientist
Broadcom Corporation, USA
Adjunct Professor, Department of Electrical and Computer Engineering
University of California

Gottfried Ungerboeck
Technical Director, Communication Systems Research
Broadcom Corporation, Switzerland

Introduction

The data over cable service interface specifications (DOCSIS) 1.0 and 1.1[1] define the physical-layer radio frequency (RF) interfaces for data-over-cable systems. With the success of DOCSIS–based broadband access to homes and businesses, the demand for greater symmetric capacity is growing rapidly. This paper discusses the physical-layer enhancements developed to increase the effectiveness of transmission in the upstream or return-path direction in which cable modems send data in time-division multiple access (TDMA) mode to the cable-modem termination system (CMTS) at the head end.

Advanced TDMA Proposal Summary

The key parameters of advanced TDMA[2] are summarized in *Table 1*.

Trade-Offs Leading to Advanced TDMA

TDMA versus S–CDMA and OFDM
Three technologies have been considered for advanced upstream transmission: advanced TDMA, synchronous code division multiple access (S–CDMA), and optical frequency division multiplexing (OFDM). For each technology, the access method during initial registration and station maintenance is TDMA. In addition, in each of the considered technologies time slots are assigned to different users, so all three schemes include TDMA burst transmission and TDMA media access control (MAC). Thus, all three schemes require a TDMA burst modem, with varying synchronization requirements. S–CDMA and OFDM can be viewed as extensions of TDMA transmission to two-dimensional (2D) framing: time/code in the case of S–CDMA and time/tone in the case of OFDM.

Impulse Noise Sensitivity: With 2D schemes, multiple codes or tones are transmitted simultaneously. Hence, individual symbols are lengthened in time by a factor of the number of codes or tones. The longer symbol duration provides an advantage in the presence of weak impulse noise, because the impulse energy is spread among the concurrently transmitted symbols. However, this advantage turns into a disadvantage if the impulse energy exceeds a certain threshold. Then more symbols are affected in these schemes than in pure TDMA.

Figure 1 illustrates the corruption of a single TDMA symbol (or a few such symbols) by a short, but strong, burst of noise, whereas the same noise event affects all concurrently transmitted CDMA or OFDM symbols.

Narrowband Ingress Sensitivity: With TDMA, agility in modulation bandwidth and carrier frequency can be used to avoid RF bands with known severe interference. Remaining narrowband interference can be mitigated by adaptive ingress cancellation. With S–CDMA, cancellation technique can in principle be employed at the chip rate. However, the large decision delay due to the inherent chip-to-symbol conversion process prevents instant availability of reliable decisions. Other estimation and subtraction techniques of much greater complexity must be employed[6,7].

TABLE 1

Key Parameters of Advanced TDMA Physical Layer

Category	DOCSIS 1.0/1.1	Advanced TDMA
Modulation	QPSK, 16-QAM	QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM
Symbol rates (Mbaud)	0.16, 0.32, 0.64, 1.28, 2.56	0.16, 0.32, 0.64, 1.28, 2.56, 5.12
Bit rates (Mbps)	0.32 –- 10.24	0.32 – 30.72
FEC	Reed-Solomon, T = 0 to 10	Reed-Solomon, T = 0 to 16
Interleaving	None	RS byte; block length may be adjusted dynamically to equalize interleaving depths
Equalization	Transmit equalizer with 8 T-spaced taps	Transmit equalizer with 24 T-spaced taps
Ingress mitigation	Vendor specific	Receiver ingress cancellation
Preamble	QPSK or 16-QAM; length ≤ 1024 bits	QPSK-0 (normal power) and QPSK-1 (high power); length ≤ 1536 bits (≤ 768 T)
Spurious emissions	Sufficient for 16-QAM	Generally 6 dB tighter to support 64-QAM

TABLE 2

Cable-Plant Impairments and Mitigation Techniques

Impairment	Mitigation Techniques
AWGN	Minimize implementation loss (see *Figure 3*)
Impulse/burst noise	FEC and interleaving (see performance section)
Narrowband ingress	Ingress cancellation (see *Figure 5*), frequency avoidance
Common path distortion[3]	Same as ingress
Micro-reflections	Transmit equalization
Hum modulation	Receiver tracking loops

To cope with dynamically occurring narrowband ingress, S–CDMA can use spreading as a form of signal repetition, trading bandwidth efficiency for robustness against narrowband noise. For example, spreading by a factor of 10 leads to a 10 decibel (dB) processing gain at the expense of a 90 percent loss of capacity—an unattractive trade-off.

In an OFDM system, narrowband interference can be avoided by not using the affected tones. Frequent tone reassignment is then required for both cable modems and the cable-modem termination systems (CMTS) to avoid dynamic ingress noise, making the protocol complicated and less efficient. OFDM is also more sensitive to unrecognized low-level narrowband interference than TDMA.

Synchronization Sensitivity: The 2D schemes require much tighter timing synchronization and lower carrier phase noise/frequency offset to maintain code or tone orthogonality. For instance, the synchronization accuracy requirement for uncoded 64-QAM is +/– 3 ns for S–CDMA, compared to +/– 250 ns for TDMA, a factor of more than 80. Thus, the timestamp-messaging scheme currently used to control DOCSIS TDMA slot timing would have to be modified. To prevent degradation of code orthogonality (the tendency of codes to interfere with one another if not perfectly synchronized and equalized), the S–CDMA system requires preshaping to flatten the band. Achieving sufficient flatness to support high-order constellation formats can be a challenge in the presence of equalization and ingress cancellation as well as imperfect synchronization among many modems spread over the cable plant. The result can be a self-inflicted noise floor due to inter-code interference.

MAC Impacts of 2D Schemes: Complexity and compatibility with the MAC and transmission convergence (TC) layers of DOCSIS 1.0/1.1 represent major hurdles for the integration of a 2D scheme. Bandwidth allocation is more complex because the scheduler must be able to schedule in time for

FIGURE 1

Strong Impulse Noise (Vertical Black Band) Affects a Larger Number of Symbols in 2D Schemes

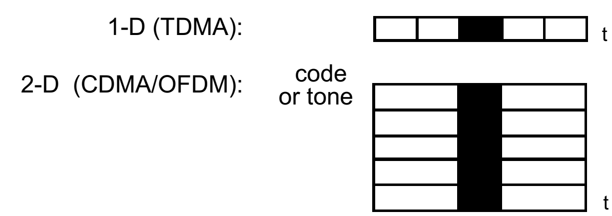

DOCSIS 1.0/1.1 modems and time/code or time/tone for advanced physical-layer (PHY) modems. This requires a change to the MAC/PHY interfaces and tight coupling between the MAC and PHY. Otherwise, there may be unused codes/tones, causing capacity loss. Because 2D schemes require block-based processing based on frames, the latency is increased, which can pose a problem for delay-sensitive applications such as voice over IP (VoIP). Some versions of OFDM utilize dynamic bit loading per tone, requiring modifications to the fixed mini-slot format (i.e., constant bytes per time interval) in the DOCSIS MAC.

Fidelity Requirements versus Implementation Complexity

An important element of advanced TDMA is its tighter specifications on spurious emissions. To a first approximation, with 64-QAM operation on the plant, cable modems must provide 6 dB higher suppression of out-of-band emissions than with DOCSIS 1.0/1.1, which uses at most 16-QAM transmission. Otherwise self-noise could limit system performance. The spurious emissions requirements were tailored to the performance achievable with practical low-cost power amplifiers. Similarly, carrier phase noise and transmitter modulation error ratio (MER) were tightened to preserve performance for 64-QAM while weighing implementation complexity and the availability of low-cost oscillators.

System Implementation

Figure 2 illustrates a typical upstream transmitter.

Figure 3 shows a typical burst receiver. The receive equalizer may include an adaptive ingress canceller. The analog front end may include a high–sampling-rate analog-to-digital converter (ADC), which allows direct digital sampling of the entire upstream band (5-42 or 5-65 MHz). Channel-quality–monitoring capability utilizing an FFT can be integrated into the receiver to support spectrum management and provide the ability to detect, characterize, and avoid interference[5].

FIGURE 2

Block Diagram of an Upstream Transmitter

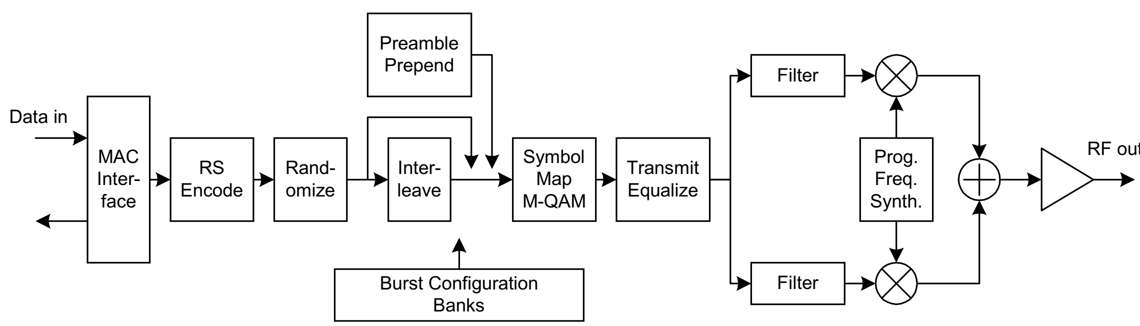

FIGURE 3

Block Diagram of an Upstream Burst Receiver

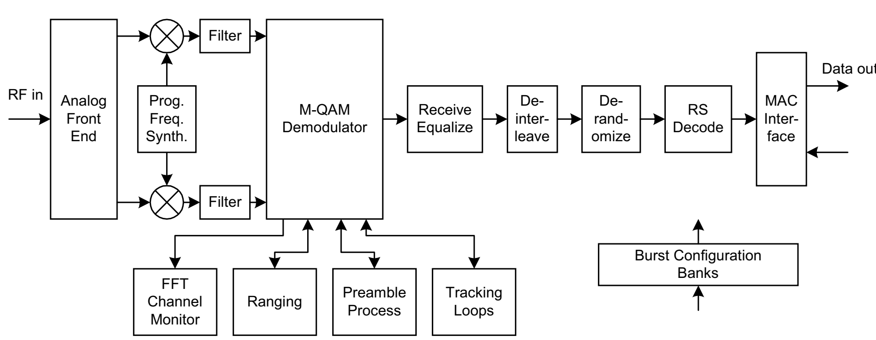

Performance

Measured Packet Error-Rate Performance

Figure 4 shows measured performance of an advanced TDMA burst receiver in AWGN with 64-QAM modulation and short packets RS(99,73). The theoretical curve for an ideal continuous receiver with RS(99,73) forward error correction (FEC) is shown for reference. The performance reflects the use of modern digital receiver design techniques and a high level of very–large-scale integration (VLSI).

Measured Ingress Mitigation Performance

Figure 5 displays measured samples of a 64-QAM Advanced TDMA signal with four narrowband interferers. The receiver ingress canceller removes the interference.

Impulse/Burst Noise Performance

The advanced TDMA physical layer adds interleaving as an important mitigation feature against burst noise, used in combination with FEC. DOCSIS contains a great deal of flexibility in burst parameters such as modulation, preamble, and FEC, which permit a trade-off of receiver processing complexity versus robustness. In this section, we consider only interleaving and FEC, since they are defined by the transmission waveform and provide a performance bound.

Specification and Measurement of Burst Noise: Burst noise is characterized statistically by its level, duration, and interarrival time.[4] Some sources of burst noise are quasi-periodic—for example, AC-line–based noise. Performance is often specified and measured in practice using a periodic burst noise model.[5] The burst noise level can reach C/I = –10 dB in the signal bandwidth. The duration is normally 1 us or less, with infrequent isolated occurrences up to 50 us, and an average interarrival time of about 10 ms (100 Hz repetition rate). System performance is measured by the packet error rate (PER) versus FEC code rate (k/n in the RS codewords). To characterize performance in the presence of burst noise, we ignore fragmentation and concatenation and settle on two payload sizes: 74 bytes (for a typical short packet) and 1528 bytes (for a long Ethernet packet).

Performance: *Table 3* shows the performance of advanced TDMA interleaving/FEC in the presence of burst noise for Reed-Solomon T = 16. The analysis assumes the corruption of all symbols coincident with the noise burst plus an additional symbol before and after the noise burst.

Figure 6 is a graphical presentation of the data in *Table 3*. Each point is a bound on the burst noise duration and repetition rate that can be corrected by the selected settings for the interleaver, FEC, and modulation. Hence, any duration and repetition rate of burst noise less than these values can be corrected. This results in a rectangular region to the left and below each point. Case "E" is selected as an example in the graph. The region of typical burst noise that occurs in real cable plants is also depicted; considerable margin is available in most cases. Actual performance, including acquisition effects, will approach these bounds but should lie below them.

FIGURE 4

Measured 64–QAM Packet Error-Rate Performance for Short Packets

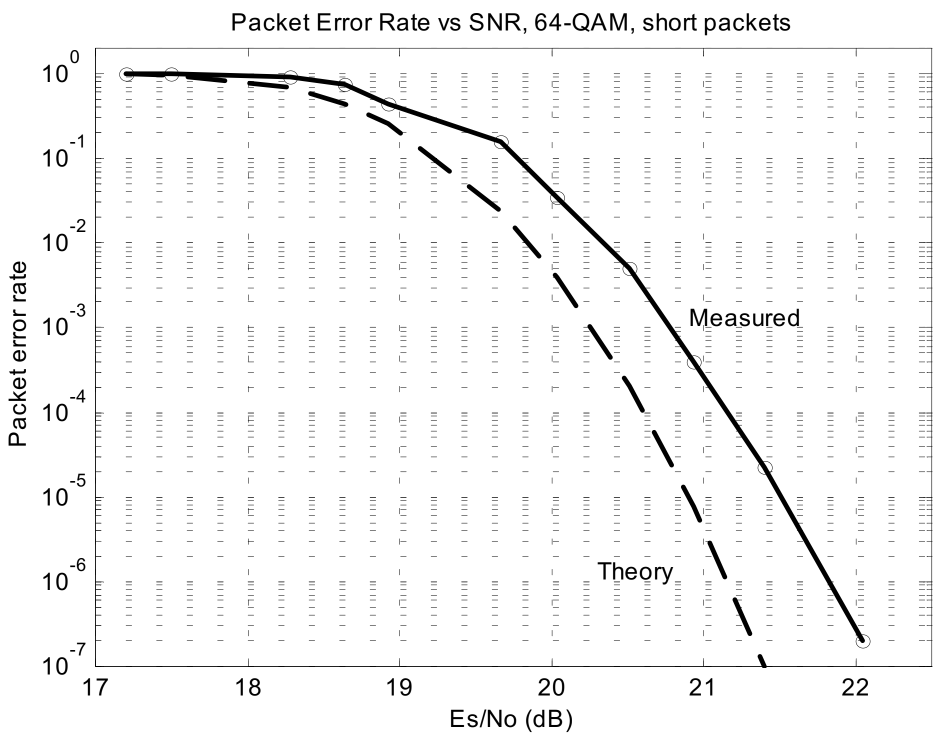

FIGURE 5

Cancellation of Narrowband Ingress. (a) FFT monitor display showing 64–QAM signal with four narrowband interferers. (b) Signal constellation at slicer input before ingress cancellation is enabled. (c) Signal constellation with ingress cancellation enabled.

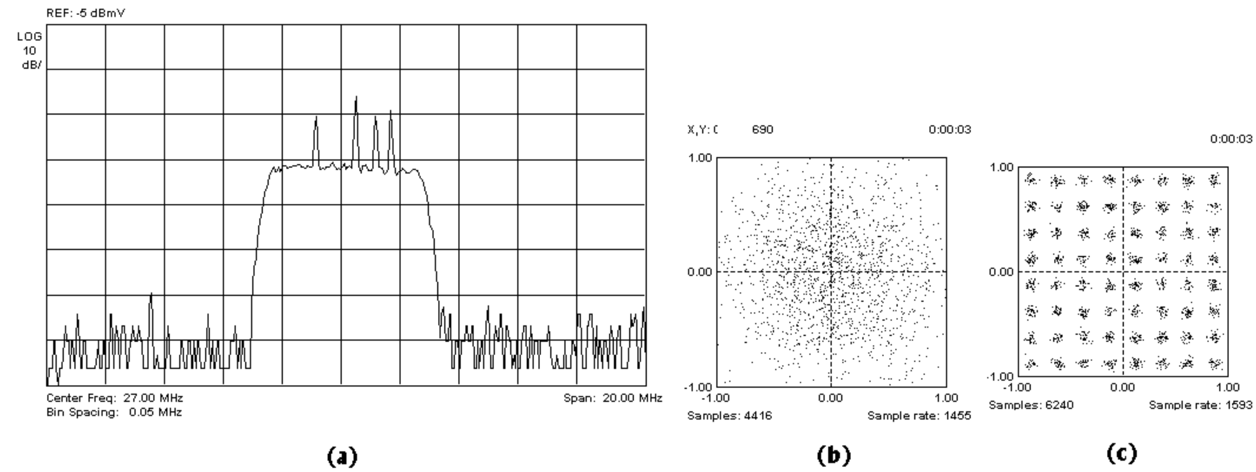

Conclusion

The proposed enhancements to the DOCSIS physical layer for upstream transmission provide increased capacity in the return path of data-over-cable systems. The approach is to extend the successful DOCSIS system in an evolutionary manner for both higher throughput and increased robustness. Attention is given to backward compatibility with 1.0/1.1 modems, coexistence between advanced TDMA and legacy modems, and specifications to permit interoperability between multiple vendors that will implement the enhanced waveform. Advanced TDMA modems employ digital techniques to mitigate worst-case plant impairments. This permits the use of portions of the upstream band that were previously unavailable due to strong impulse noise or narrowband ingress. The result is expanded capacity of data-over-cable systems in support of new broadband services demanding higher symmetrical throughput, including voice, video telephony, videoconferencing, and distributed servers.

TABLE 3

Interleaver/FEC Performance in Periodic Burst Noise

Case Label	Modulation Format	Symbol Rate (Msps)	Packet Length (bytes)	Number of Interleaved RS Codewords	Maximum Correctable Noise Burst Length (usec)	Maximum Correctable Noise Burst Repetition Rate (kHz)	FEC Code Rate (%)
A	64-QAM	5.12	74	No interleaving	3.3	35.7	69.8
B	64-QAM	5.12	74	2	7.4	27.7	53.6
C	64-QAM	5.12	74	4	15.2	18.9	22.4
D	64-QAM	5.12	1528	32	64	3.01	59.9
E	64-QAM	5.12	1528	64	131	2.15	42.7
F	16-QAM	1.28	74	No interleaving	21	6.06	69.8
G	16-QAM	1.28	74	2	44	4.69	53.6
H	16-QAM	1.28	74	4	93	3.19	22.4
I	16-QAM	1.28	1528	32	390	0.502	59.9
J	16-QAM	1.28	1528	64	790	0.358	42.7

FIGURE 6

Interleaver/FEC Performance in Periodic Burst Noise

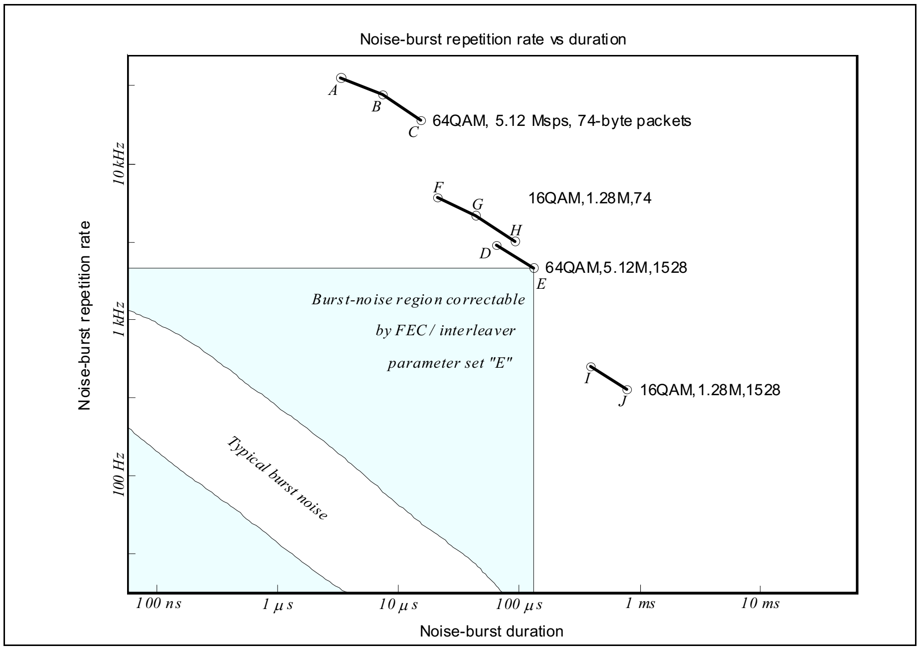

Acknowledgement

The authors would like to acknowledge the dedication of their colleagues who participated in the trade-off, specification, design, fabrication, characterization, and test of the advanced PHY system.

Notes

1. Data-over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SP-RFIv1.1-I05-000714, Fifth Interim Release, July 14, 2000.
2. *Advanced TDMA Proposal for HFC Upstream Transmission*, Broadcom Corporation and Texas Instruments Cable Broadband Communications, submitted to CableLabs in December 1999.
3. For channel model information on plant impairments see in *Cable Modems: Current Technologies and Applications*, International Engineering Consortium, Chicago, 1999, the following papers:
   - Thomas J. Kolze, “An Approach to Upstream HFC Channel Modeling and Physical-Layer Design”
   - Bruce Currivan, “Cable Modem Physical Layer Specification and Design.”
4. R. Prodan et al., “Analysis of Two-Way Cable System Transient Impairments,” CableLabs, NCTA Proceedings, 1996.
5. Daniel Howard, “System Identification and Adaptation to RF Impairments Using Advanced PHY,” co-publication in this conference proceedings.
6. Lops, Ricci, and Tulino, “Narrow-Band-Interference Suppression in Multiuser CDMA Systems,” *IEEE Transactions on Communications*, Vol. 46, No. 9, September 1998.
7. Young and Lehnert, “Analysis of DFT-Based Frequency Excision Algorithms for Direct-Sequence Spread-Spectrum Communications,” *IEEE Transactions on Communications*, Vol. 46, No. 8, August 1998.

Technology Roadmaps for the Optical Internet Data Center

Casimer DeCusatis
Senior Engineer
IBM Server Division

Larry Jacobowitz
Photonics Manager
IBM Microelectronics Division

Abstract

In this paper, we review some of the technology trends and directions in fiber-optic data communications, including wavelength multiplexing and storage networking, and discuss applications of this technology, such as disaster recovery solutions and grid computing. We identify key hardware components that enable next-generation data centers, including wavelength agile lasers, optical switching, and parallel optical links. Related trends in self-managed computer networks and emerging industry standards such as InfiniBand are also presented.

1. Introduction

In the past decade, the computer and data-communications industries have witnessed far-reaching changes in both the architecture of large, enterprise-class computing and the fiber-optic networks that are used to interconnect servers with clients, storage, and each other. With the widespread adoption of fiber-optic technology, network bandwidth may have surpassed even processor speed as the most critical element in evaluating the utility of large computer systems. As modern data centers have come to rely on the Internet, which in turn has come to rely on optical communications, various technologies and applications have emerged that are expected to drive future growth in these areas. In this paper, we describe a few of the most important trends in this field, including wavelength multiplexing for metropolitan-area networks (MANs), storage networking, grid computing, automated network management, the emerging InfiniBand standard, and others.

2. Wavelength Division Multiplexing

Wavelength division multiplexing (WDM) is emerging as the preferred technology for the optical data center of the future. It takes advantage of the fact that different wavelengths or colors of light will not interfere with each other when they are carried over the same physical optical fiber. The concept is similar to frequency multiplexing used by frequency modulation (FM) radio, except that the carrier "frequencies" are in the optical portion of the spectrum (around 1550 nanometers wavelength, or 2 x 10^{14} hertz). Thus, by placing each data channel on a different wavelength (frequency) of light, it is possible to send many channels of data over a common optical fiber [1]. Using optical amplifiers, DWDM networks can be extended to well over 100 km distances, making them well suited for disaster recovery applications. More data channels can be carried per fiber if the wavelengths are spaced closer together. In this manner, WDM systems may be classified as coarse, wide-spectrum, or dense WDM (DWDM), as shown in *Table 1*. Successive generations of DWDM equipment have supported more wavelengths and thus more channels over a common fiber link. Current industry standards ratified by the International Telecommunications Union (ITU) establish a minimum spacing of 0.8 nm (100 GHz) and accommodate about 32 wavelengths per fiber. Each wavelength today has a typical capacity of 2.5 gigabits per second (Gbps), which is expected to increase to 10 Gbps in the near future. Systems are emerging that are expected to handle hundreds or even thousands of wavelengths at less than half the current wavelength spacing.

The next generation of DWDM involves more than simply increasing the number of channels, however. New network topologies are also being enabled by the capability of third-generation and higher DWDM systems to support meshed or nested rings with dual redundant paths, self- healing at the physical layer, data regeneration, and more advanced survivability or path protection than currently available. By contrast, conventional telecommunications networks have deployed asynchronous transfer mode (ATM) over synchronous optical network (SONET) rings with separate overlay networks to accommodate Internet protocol (IP) data and other kinds of traffic. These overlay networks are optically transparent but remain service specific. Since data traffic has different characteristics than voice traffic, overlay networks cannot make efficient use of the available bandwidth as they scale to multiple, concatenated rings. Consequently, optically transparent overlay networks are being replaced by a service-transparent DWDM core capable of allocating bandwidth on demand. This offers the advantages of a highly scalable, low-cost, protocol-independent infrastructure and may be the first step toward all-optical networks (AONs). Unlike first-generation DWDM, more recent technology allows wavelength multiplexing networks to be cascaded together, because they act as a complete "3R" repeater (Retimes the signal to remove jitter and improve clock/data

TABLE 1

Types of Wavelength Multiplexing Systems

Type of WDM	*Number of Wavelengths (Channels)*	*Wavelength Spacing*	*Notes*
Coarse (CWDM)	2–4	Very wide (typically 1300 nm and 1550 nm)	Very limited applications
Wide Spectrum (WWDM) also known as Sparse (SWDM)	Up to 16	No standard defined; typical spacing of 1 to 30 nm have been used	Four-channel systems are under consideration for emerging 10 Gbps Ethernet[1]
First-Generation DWDM	8 – 10	ITU grid[2]	Example: IBM 9729
Second-Generation DWDM	Up to 16	ITU grid	
Third-Generation DWDM	Up to 32	ITU grid	Largest systems commercially available for datacom Example: IBM 2029
Fourth-Generation DWDM (Ultra-Dense)	Up to 64	No standard defined; smaller than ITU grid (0.4 nm proposed)[3]	Not commercially available

Notes:
1. See IEEE 802.az higher-speed discussion group (10 Gigabit/second Ethernet) at http://grouper.ieee.org/groups/802/3/10G_study/public/
2. ITU grid nominally centered on 1550 nm with minimum wavelength spacing of 0.8 nm (100 GHz) or multiples thereof, anchored to a reference of 193.1 THz per ITU G.MCS Annex A of COM15-R 67-E.
3. See, for example, M. Ferries, "Recent Developments in Passive Components and Modules for Future Optical Communication Systems", paper MI1, Proceedings of the. OSA Annual Meeting (Santa Clara, Calif.: 1999), 61.

(See C. DeCusatis, "Dense Wavelength Division Multiplexing for Parallel Sysplex and Metropolitan/Storage-Area Networks," Optical Networks, Vol. 2, No. 1 (January/February 2001), 69–80.)

recovery, Reshapes the signal to removes pulse distortion caused by dispersion, and Regenerates the signal to insures that there is sufficient optical power to reach its destination). Full protection switching at the physical or transport layer is available on a per-channel basis to restore service in the event of either a fiber break or failure of a hardware component in the system; a properly designed DWDM network would have no single points of failure [2].

The historical trend of growth in aggregate system bandwidth is likely to continue for the near future. However, this approach will require the service layer to upgrade more than twice as fast as the transport layer, or roughly double capacity every six months (some estimates have shown the service layer growing more than 70 times by 2003). A more realistic approach is to have the transport and service layers evolve together, although this still requires service-layer capacity to double on a yearly basis. To keep pace with bandwidth growth, future DWDM systems may employ ultra-dense fourth-generation DWDM systems. The desire for efficient bandwidth management is also likely to drive the use of sub-rate multiplexing or the combination of multiple time division multiplexed data channels within a single wavelength on a DWDM network as the most cost-effective way to increase the total number of channels in a DWDM network. As the capacity of the network grows, there are a number of key DWDM components that must keep pace. For example, this past year saw the introduction of the first DWDM devices to employ both electrical and optical backplanes as an efficient means to handle the huge volumes of data being routed within networking equipment. These optical backplanes are also being extended to adjacent equipment racks using parallel optical interconnects, or ribbons of optical fiber four to 12 channels wide, capable of supporting up to 2.5 Gbps/channel. Related technologies, such as all-optical switching and routing, are also being developed to keep pace with bandwidth demands in the network. As the number of wavelengths is increased, there is also a need for optical sources that can be quickly tuned to different wavelength bands to provide more flexible optical add/drop capability and reduce the number of components required for servicing the network. Currently, the mismatch between laser center wavelengths and optical filters in a DWDM network can result in 3 dB to 4 dB or more of excess attenuation, effectively limiting the supported distance that can be achieved without using optical amplifiers. Tunable laser sources would also alleviate this problem. The current state of the art for tunable lasers [3] involves changing the laser wavelength by tightly controlling its temperature, often using some form of feedback to a thermoelectric cooler or a specially designed laser structure. This approach is commonly known as wavelength locking. Currently available systems exhibit fairly slow tuning times, stability concerns over long periods of operation, a limited tuning range, and high cost. Future advances in this area are expected to be a key enabler in lowering the cost and increasing the flexibility of DWDM networks.

Previously, features such as quality of service (QoS), guaranteed message delivery, protection switching, and others have been provided over the legacy telecommunications infrastructure. Indeed, SONET–based traffic is already carried over a physical layer that uses DWDM optical fiber interfaces. However, the growth of IP traffic has led to a complicated arrangement with up to four separate transport layers (IP over ATM over SONET over DWDM). In an effort to simplify this approach and streamline the data flow, there is a clearly emerging trend toward elimination of the ATM and SONET layers and toward the direct transmission of IP or grid data over DWDM. This new model requires the DWDM layer to assume many of the traditional functions associated with ATM over SONET, such as protection switching. In particular, two standards efforts are currently under way to link data packets directly to DWDM optical wavelengths so that the optical network can take some advantage of the intelligence imbedded in IP traffic. This so-called "optical IP" effort could eventually allow grid users to dynamically request portions of a fiber cable's bandwidth for a particular time or service. One such standards effort is the Optical Domain Service Interconnect (ODSI) coalition—a loose connection of vendors providing optical transmission equipment, access services, terabit routers, switches, and network provisioning software. ODSI seeks to define common control interfaces between optical or electrical physical layers and IP media access layers of the open systems interconnect (OSI) model. This work may ultimately rely on derivatives of the Internet Engineering Task Force (IETF) multiprotocol label switching (MPLS)—a means of defining IP flows that is already widely used in the electrical signaling domain. ODSI will propose low-level (below Layer 3) control plane standards that must be met by vendors of both optical transmission equipment and broadband IP switches/routers. A separate but complimentary effort, also based on MPLS, is being drafted within the IETF itself. Known as multiprotocol lambda switching (MP_S)—lambda refers to switching by native wavelengths—this effort has proposed methods to link optical cross-connects and high-bandwidth routers through a Layer-3 switching methodology. Operating at a higher level than ODSI, this proposal has the advantage of not requiring a new set of protocols for QoS and bandwidth control

3. Storage-Area Networking

Today's strategic business initiatives are inexorably linked to network bandwidth for storage applications. Simply put, a business can only grow as fast as information can be retrieved, exchanged, and acted upon. A recent University of California study [4], having estimated that the entire human race has accumulated about one exabyte (10^{18}) of information to date, goes on to predict that the second exabyte will be generated within the next three years. This abrupt growth spurt in storage requirements after a fairly long period of incremental growth parallels the growth pattern of the Internet, and its implications are just as fundamental. It has been estimated that by 2003 the Fortune 1000 companies will add more than 150 terabytes of storage capacity and that more than half of all fiber-optic traffic in their data centers will be feeding storage systems. Storage services are growing at more than 120 percent per year, and are projected to be an $8 billion market in 2003 [5]. Loss or even temporary unavailability of mission-critical data is unacceptable. In the event of a disaster at their primary data center, large companies can easily lose millions of dollars. In this environment, storage-area networks (SANs) have become a vital part of the e-business infrastructure.

With the growing amounts of network bandwidth that must be dedicated to storage applications, it makes sense to separate the storage traffic from the rest of the network to improve performance and simplify network management. This is the basic principle behind a SAN. The concept of SANs is not new. More than 10 years ago, enterprise data centers used Enterprise Systems Connectivity (ESCON™) protocols over a dedicated, switched fiber infrastructure as a replacement for copper cables. Optical networking quickly became the preferred means to interconnect mainframe servers and storage devices. More recently, Fibre Channel protocol (FCP) has become the preferred means of building a SAN because it offers higher bandwidth and improved performance at extended distances (beyond 10 km) for disaster recovery applications. This is the technology currently being deployed by many service providers for storage over a wide area and metropolitan area. Examples of this approach include the recently announced GeoMax℠ service from Qwest and the Ultravailable Broadband Network℠ from AT&T, both of which are based on DWDM technology.

In addition to offloading storage traffic and minimizing server-to-server traffic to improve network performance, SANs are typically optimized for large blocks of data transfer (streaming video or terabyte back-ups with many sequential read/write operations). SAN management is centralized (thought it may not be done from the data server), and many suites of storage-management tools are available. Storage resources can be dynamically re-allocated and shared across multiple servers in a SAN. This allows heterogeneous SANs to be constructed that are still able to take advantage of policy-based management, workload balancing, and similar functions (the cost of network management can be 3 to 10 times greater than the SAN hardware cost). An important alternative to FCP–based SANs is the Fibre Connection (FICON™) protocol [1], originally developed for enterprise servers and currently being standardized by the American National Standards Institute (ANSI). FICON is implemented in essentially the same way as FCP, below Layer-3 protocols, and includes enhancements above Layer 4 that improve data integrity. Through various design features, FICON channels maintain their maximum data throughput over distances of at least 100 km.

A complimentary approach to SANs known as network attached storage (NAS) relies on IP protocols such as Gigabit Ethernet to interconnect application servers with storage appliances. NAS does not segregate storage traffic from other types of IP data and is typically optimized for smaller data packet transfers, using file protocols rather than block protocols (the maximum data block size for Gigabit Ethernet is 9 kilobytes using jumbo frames). The principle benefit of NAS is its flexibility and ease of installation. Storage appliances are designed to "plug and play" on existing local-area networks (LANs) with a minimum of effort. The difference in complexity is apparent; while you install a NAS device, you implement a SAN infrastructure. Thus, NAS architectures can vary from shared nothing (each server accesses storage independently with no clustering, load balancing, or failover capabilities) to shared everything (any server can concurrently access any storage device). Various types of SAN/NAS gateways are available to merge the two approaches and groom storage data traffic, for example the IBM TotalStorage product family.

Storage applications over IP are generating interest because of their ability to use existing IP network infrastructures over MAN or wide-area network (WAN) distances. Although transport of Gigabit Ethernet over DWDM has been demonstrated to distances greater than 1,000 km [3], performance at such long latencies remains to be determined. There are currently three major ways to transport storage data over IP. First, iSCSI (SCSI over IP) uses a software agent on the server to encapsulate SCSI data into an IP packet, and then tunnel that encapsulated packet through the transmission control protocol (TCP)/IP stack. The data is unwrapped at the other end of the link. Second, storage over IP (SoIP) converts Fibre Channel to IP outside the server and bypasses the TCP/IP stack in the host, eliminating stack overhead in the process. Finally, the Fibre Channel over IP (FCoIP) specification is a proposed standard that allows routers to connect a Fibre Channel SAN to an IP network (note that similar options for Fibre Channel over ATM are also available). Of these options, iSCSI has attracted attention recently because of its handling of block input/output (I/O) protocols and ready integration with embedded storage (iSCSI appliances) or gateways to an existing SAN or NAS. As a block transfer–based network accessible over an IP network, iSCSI has the potential to offer the advantages of both SAN and NAS environments.

The fundamental building blocks of a SAN include the server, optical transport fabric (gateways, hubs, wavelength multiplexers), switches or routers, storage devices (disk control units or tape libraries), management software, and deployment service offerings. Steadily increasing server performance has driven the requirement for increased network bandwidth and storage capacity. In the past decade alone, the cost per megabyte of disk storage has fallen by an order of magnitude to less than 50 cents. Data density on a storage medium has been doubling annually since 1997, growing faster than Moore's Law. Although optical storage is becoming more common, magnetic storage continues to grow due to recent advances in material science [6] that are expected to enable 100 gigabits/square inch by 2003—a level previously thought impossible (this would enable a laptop PC to hold 200 gigabytes, enough to accommodate 200,000 books or 300 CDs). There is also a significant opportunity to improve the optical transport fabric technology. For example, DWDM has emerged as a cost-effective way to extend storage networks over the MAN and WAN and is expected to form the backbone for future storage service provider offerings. SONET backbones will likely co-exist with Ethernet encapsulation for many years to come, especially in the MAN. This in turn has driven interest in very high-speed (nanosecond) all-optical switches and cross-connects. These technologies are among the critical elements of the future optical data center.

Today's storage environment is a combination of multiple SAN and NAS environments, typically with little attempt at global optimization. In the future, storage networks will evolve through the use of gateways, modular storage, and hybrid approaches such as iSCSI into a shared pool of intelligent storage devices accessible from any server, anywhere. Efforts are under way toward storage virtualization, which separates a single, logical view of storage resources from the physical device configuration. However, a standardized approach has yet to be established. Intelligent storage devices are expected to automate storage administration, making it simpler, less expensive, and continuously available. With the growth of DWDM as a cost-effective means to implement disaster recovery over extended distances, storage networks are spreading into the MAN and WAN. Given the recent surge of interest in business continuity solutions, this trend toward geographically dispersed SANs is expected to continue and represents a significant market for storage application providers.

4. Grid Computing

Grid computing involves many thousands or millions of small, distributed computing devices, linked together in local or regional networks that are in turn interconnected on a global level. The grid functions as if it were one giant virtual computer, with very close integration of servers, storage, and other resources. Everyone connected to the grid is able to share these resources in a very fast, efficient manner. Grids can be used to address so-called "Grand Challenge" problems, including advanced genetics research, modeling global weather patterns, or global air traffic control. This class of high-risk/high-reward problems is also known as "Deep Computing." While they hold the potential to be the next disruptive technology in the computing industry, the concept of grid computing isn't new. It can be traced back to early experiments with distributed, parallel processing, and so-called "lightweight grids" such as SETI@home, which uses free software downloaded over the Internet to scavenge a computer's spare processing cycles to analyze signals from the Arecibo Radio Telescope. There are currently more than 1.6 million SETI@home subscribers in 224 countries, averaging 10 trillion operations per second and having contributed the equivalent of more than 165,000 years of computing time to this project. This is arguably the world's largest distributed supercomputer, interconnected over the existing Internet. Small-scale grid models have been under investigation for years by the University of Southern California, Argonne National Laboratory, and the National Aeronautics and Space Administration (NASA), among others, and commercial applications have recently begun to attract the attention of major corporations [7]. Grids are still in their infancy; however, the ultimate goal is so-called "heavyweight grids"—much larger and more powerful systems linked on a national or international scale. Today, heavyweight grids are under development in several countries, including the United Kingdom's National Grid and The Netherlands. Recently, IBM was selected to build the Distributed Terascale Facility (DTF), which would be the world's most powerful computing grid capable of 13.6 trillion calculations per second and a storage capacity of more than 600 terabytes of data, or the equivalent of 146 million full-length novels. The Terascale Facility, with $53 million in funding from the National Science Foundation (NSF), is a joint undertaking of the National Center for Supercomputing Applications (NSCA), the San Diego Supercomputing Center (SDSC), Argonne National Laboratory, and the California Institute of Technology. Including some of the world's fastest supercomputers and high-resolution visualization environments, this grid will enable thousands of scientists around the country to share computing resources over the world's fastest research network in search of breakthroughs in life sciences, climate modeling, and other critical disciplines.

These systems are based mainly on protocols, standards, and software tools under development by Globus, an open source community led by Dr. Carl Kesselman at the University of Southern California Information Sciences Institute and Dr. Ian Foster of Argonne National Laboratory and the University of Chicago. In the same way that the Linux™ community has become a major part of open standards, Globus is working with the grid movement to help various standards and technologies to reach maturity; this includes tools to enable remote sharing of massive computing and storage resources, sophisticated resource management, scheduling and scalability routines, privacy and security tools, and other software functionality. Emerging software standards such as UDDI (universal description, discovery, and integration) are important to this effort, as well as natural languages to analyze unstructured data such as Web pages, hypertext markup language (HTML), video, and MP3 audio files. Utility pricing, or pay-as-you-go, is one of the features that make grid computing attractive to large businesses; the increased flexibility that this model affords could save billions of dollars alone. This approach is different from capacity-on-demand models, in which companies own or lease servers and pay an additional, fixed charge when they activate excess processing capacity. By contrast, in the utility model, the company doesn't own anything, but rather simply pays for its use of servers, storage, and disk capacity. These resources are metered through software measurement tools and can be billed in various ways, usually either on a per-user or per-processor basis. Either way, the cost per user goes down as the server utilization rises. This approach avoids one of the major pitfalls of today's information technology (IT) systems, namely capacity planning. Many corporations have difficulty accurately forecasting their demand for IT resources, and this has become more obvious during the recent economic slowdown, when anticipated revenues fail to materialize and overbuilt IT capacity becomes a burden that most companies would prefer not to bear. Since utility pricing directly ties cost to revenue, it should allow companies to weather fluctuations in the market better by amortizing costs over time.

5. Autonomic Computing

The cost of ownership for servers is already being driven upward by management and maintenance fees. In fact, some feel that the management of ever-growing networks constitutes a looming crisis in the Internet data center. The rapid growth of SANs, DWDM networks, and grid computers has only accelerated this trend. For these systems to succeed, both resource and network management must be highly automated. The future of optical network management depends on a core of intelligent middleware that is capable of handling most networking decisions without human intervention. This is the concept of autonomic computing, named after the body's autonomic nervous system that allows human beings to adapt to any number of situations by unconsciously adjusting things such as your heart rate, breathing, body temperature, or pupil dilation. An autonomic computer is self-managing and self-healing, responsive to changes in its environment, and always accessible. This is accomplished by middleware needs to be platform agnostic, supporting a heterogeneous mixture of access devices ranging from conventional PCs to network appliances, embedded devices, kiosks, or cell phones. In addition to network management, middleware may also perform functions such as data caching to improve performance, directory and security control, QoS enforcement, and transcoding (changing data structures to take on the characteristics of the end user's access device).

A leading example in the field of self-managed computer systems is IBM's project eLiza, announced in April 2001 [8]. This refers to the first artificial intelligence program to permit natural language conversations between a human and a computer, written decades ago by Professor Joseph Wiezenbaum. The symbol for the eLiza program is a lizard, which refers to calculations done a few years ago using assumptions in Ray Kurzweil's book "The Age of Spiritual Machines," in which researchers estimated the processing power of the chess-playing Deep Blue supercomputer as roughly equivalent to a lizard's brain. While it may seem odd for such a simple creature to need so much brainpower, the lizard actually has a far more complex task than a single purpose computer. While the computer's environment is completely controlled and comparatively predictable, the lizard faces a constantly changing jungle environment requiring instantaneous responses to new demands and dangers. Enterprise computers have their own jungles to contend with, complete with totally unpredictable demands from unknown amounts of users, sudden threats from predators, and frequent natural shocks that stress the system. Like the lizard, computers require intelligent self-management skills to survive in this environment. A self-managed computer could, for example, update and maintain its own software at the latest release levels, actively protect itself against attacks by malicious computer hackers, monitor its health and perform limited acts of self-repair, correct inadvertent human management errors, and guard against unexpected problems or acts of nature. In this way, an eLiza system would insure its own survival and stability. It would also continually self-optimize, eliminating the need for constant tweaking by network administrators, and adjust by itself to changes in the server loading. For example, an eLiza server running the stock market could detect sharp upsurges in trading and adjust quickly to meet the higher demand. In addition to providing the means to control systems that are orders of magnitude larger than any existing today, it is hoped that these systems would be significantly lower in cost to own and maintain. While eLiza is still under development, some important related steps have already been taken—for example, the z/OS™ 64-bit operating system automatically shifts server capacity between logical partitions in response to high-priority workload spikes. This is expected to be a great productivity boost for IT staff, freeing them from the more mundane tasks to focus on other areas. It may also help alleviate projected shortages in trained IT staff. (Some analysts estimate that in five years, the world will be short at least a million qualified IT administrators. Others have noted that a network of one billion users would require about 250 million skilled people, almost the entire population of the United States) It is widely recognized that there are some things that computer systems cannot do for themselves, and although their job descriptions may change, systems administrators are not going to disappear. However, if eLiza succeeds, the end result could be a global computing network that will be as easy to control as a kitchen appliance.

6. The InfiniBand Standard

Features that used to be associated exclusively with mainframe computing—continuous availability, serviceability, high reliability, redundancy without single points of failure, clustering, logical partitioning, and greater I/O bandwidth—are slowly becoming commodity requirements in today's e-business environments. This had led to an increasing focus on the limitations of current system architectures, in particular widely used standards such as the peripheral component interconnect (PCI) bus. While microprocessor and memory subsystems have advanced at a rapid pace during the last decade, high-performance I/O has traditionally been reserved only for large mainframes or enterprise servers, which implemented direct channel attachment for I/O devices. In an effort to develop new industry standards in this area, the InfiniBand Trade Association (IBTA) grew from a combination of two previous industry consortiums, Next Generation I/O (NGIO) and Future I/O. In August of 1999, the IBTA was founded by a core group of companies that form its steering committee and current board of directors, namely IBM, Intel, Hewlett-Packard, Compaq, Dell, Microsoft, and Sun Microsystems. They were joined by many companies representing virtually every aspect of the computer and networking industry, including sponsoring members such as Agilent, Cisco Systems, Brocade, 3Com, Adaptec, EMC, Fujitsu, Siemens, Lucent Technologies, Nortel Networks, Hitachi, and NEC. At the end of 2001, there were 226 member companies participating in the IBTA. The initial InfiniBand Architecture Specification was released to the industry on October 24, 2000, at the second annual InfiniBand (IB) Developers Conference. Since then, there have been more than 100 product announcements for silicon chipsets, adapter cards, switches, software, cables, and test equipment. The value proposition of InfiniBand lies in a combination of performance improvement and better reliability (at least for entry-level and midrange systems) in a scalable, low-cost architecture. IB is a bottoms-up design for high-performance interconnects, suitable for many different server platforms and applications ranging from SANs to server clustering. Originally conceived as a replacement for PCI–type data buses, IB has broadened its focus to encompass client-server communications, server clustering, and storage. IB provides a message passing fabric in which data is shared between a host and target channel adapter. One objective of the IBTA was to introduce an I/O subsystem with significantly higher performance than conventional shared bus architectures, with characteristics inspired more by the enterprise server environment. This approach brings with it improvements in reliability, error detection, and fault isolation. For example, fault zones in the IB fabric insure that a failure on an IB–attached I/O device will not compromise the rest of the server, as in a shared bus architecture. Redundant IB fabrics and transparent multi-pathing may also be configured for high-availability applications. Although data delivery is still on a best-effort basis, similar to IP or Internet messaging, it also offers selectable QoS features that can be provisioned by specified service levels, automatic failover, and virtual lanes. IB performance is also more scalable than its predecessors; it can readily accommodate faster, higher-performance I/O devices, and performance of the IB subsystem is not reduced by the addition of extra storage or networking capacity. It is designed to reduce high TCP/IP stack latencies between transaction servers, database servers, and load-balanced Web servers, as well as between servers and storage in future Internet data centers. InfiniBand is the first standard to incorporate parallel optical data buses to increase bandwidth and distance beyond the capabilities of traditional copper buses. Duplex IB link speeds can scale from 500 megabytes per second to 6 gigabytes per second, per link [1]. By contrast, PCI–X offers a peak burst bandwidth of 1 gigabyte per second half duplex at 64 bits, with next-generation plans for 2 to 4 gigabytes per second still in the process of being defined. There can be thousands of subnets in an IB fabric, each in turn serving thousands of nodes (servers, storage, switches, routers, network analyzers, and other devices). In contrast to other transport protocols such as Ethernet or SCSI, InfiniBand allows for the adaptation of packet sizes and transport characteristics to a given application, such as block I/O or data streaming with large or small packet sizes. There have been a number of IBTA–hosted meeting and industry trade shows featuring demonstrations of the IB fabric running Linux-based clustered server and storage applications.

7. Conclusions

There are a number of key trends currently driving the next generation of computer networking technology for the optical Internet data center. The recent exponential growth rate of storage applications, coupled with a renewed interest in disaster recovery, has accelerated deployment of fiber-optic DWDM networks for both long-haul and metropolitan-area networks. This technology functions as a protocol-independent channel extender and allows network capacity to grow without the expense of installing new fiber. This has enabled applications that were previously not economically feasible, such as multi-terabyte and petabyte SANs or mainframe clustering in a Parallel Sysplex architecture. Optical networks are also expected to play a role in emerging grid computers—massive interconnections of servers over long distances that offer computing power on a utility-based or pay-as-you-go model. As networks scale to ever-larger dimensions, some form of autonomic management or self-healing is required to address practical problems with management overhead. These advances in networking have led to interest in new industry standards to improve performance and reliability for commodity-based computers. One very promising effort is the InfiniBand standard, which functions as a replacement for the legacy PCI data bus and also enables server clustering and storage networking. The synergistic combination of these trends is expected to create new market opportunities for service providers as well as traditional computer companies capable of providing end-to-end service and support of increasingly complex networks.

8. References

[1] C. DeCusatis (Editor), *Handbook of Fiber Optic Data Communication, 2nd Edition*, and *Fiber Optic Data Communication: Technological Trends and Advances* (New York: Academic Press, 2002).

[2] C. DeCusatis, "Dense Wavelength Division Multiplexing for Parallel Sysplex and Metropolitan/Storage-Area Networks," *Optical Networks*, Vol. 2, No. 1 (January/Febraury 2001), 69–80; see also C. DeCusatis, "Information Technology and Disaster Recovery at the World Trade Center," *Dr. Dobb's Journal* on-line (September 2001), 4.

[3] H. Shakouri, "Wavelength Lockers Make Fixed and Tunable Lasers Precise," *WDM Solutions* (January 2002), 23–32.

[4] "Future of Storage Area Networks," University of California Study, (2000).

[5] C. DeCusatis, "Optical SAN and NAS Implementations," Invited Talk, OFC 2001 (Anaheim, Calif.: March 19–21, 2001).

[6] J. Lohau et al., "Dynamic Coercivity Measurements of Antiferromagnetically Coupled Magnetic Media Layers," *Appl. Phys. Lett.*, Vol. 78 (2001), 2,748.

[7] C. DeCusatis, "Grid Computing: The Next Big Thing," *Byte* magazine on-line (March 2002).

[8] Ed Scannell, "IBM Research into Autonomic Computer Development," *InfoWorld* (March 14, 2001), http://www.infoworld.com/articles/hn/xml/01/03/13/010313hnhorn.xml; see also Ed Iwata, "IBM to Build High-Speed Network," *USA Today* (August 2, 2001), http://www.newsbytes.com/news/01/168609.html.

A Multiservice Applications Platform

Kevin Duffy
Business-Development Manager
IBM

An overriding concern in the telephony industry is the ability to develop and deploy value-added applications to ensure the loyalty of the subscriber base and to generate revenue. One way to provide this ability is through a multiservice platform that resides in the network. IBM has produced one such product, called the IBM Resource Manager. Instead of a monolithic item that would run one service when placed in the network, the Resource Manager is a set of components that can be pulled together in different ways to provide different services, depending upon the provider's needs. This paper outlines the way different services can run on this platform and will illustrate what this type of distributed architecture can accomplish within a network. Three services are discussed here that the Resource Manager enables: integrated voice command, calling/conferencing services, and operator redirect.

Integrated Voice Command

One service that can be deployed through the multiservice Resource Manager platform is integrated voice command, the ability to pick up the phone, speak a command, and have the command carried out. *Figure 1* illustrates how this service works. One obvious application for voice command is voice-activated dialing. The user shown in *Figure 1*, for example, can simply say, "Call Wally," and the application will connect him to the requested party.

This application brings value to the network in and of itself, but it is only the tip of the iceberg in terms of possible value-added voice services. The advent of new voice portals today means that a wealth of information is available from the network by using voice. And as commands move to the edge of the network, appliances will allow access to this network information. People have become very comfortable using their phones. The next step is to be able to make requests of the phone as the user in *Figure 1* does: "Tell me about my stock portfolio," or, "Tell me if my plane is on time." This architecture provides the capability to fulfill those requests.

Wireless applications have been driving development of voice-activated services and probably will continue to do so. This is especially true as concern grows about people dialing their cell phones as they are driving; indeed, legislation is in the works in some areas to regulate such cell-phone usage. Other applications of voice-activated services are as rich as one's imagination. Among them are personal or corporate dialing services, financial services, weather reports, and sports updates.

Voice command can be activated by an off-hook trigger for a wireline phone or a star code sequence in a wireless environment. Activation leads to a flat menu structure. Most voice response unit (VRU) applications today tend to be extensions of old customer-service applications that used a structured way to get information from a customer and to reply with an answer. Flattening the structure provides an intuitive interface that allows specific questions and quick responses. This interface provides a common look and feel for all services and thus facilitates the introduction of new services.

Voice-activated services were tried five or six years ago with limited success. Why should the situation be any different today? The major difference is that the technology has come a long way. Voice-recognition software is now speaker independent, whereas in the past the name had to be recorded and recognized. In that case, if a dog barked in the background during the recording, it had to bark for every call or the software would not recognize the name.

The current implementation of voice-activated services is through a cellular code division multiple access (CDMA) network. It can synthesize speech, so it does not require programming. The data interface is self-provisioning through the Web—any personal digital assistant (PDA) or similar device can upload it. It can then be maintained through the Web, a process with which most people are now very comfortable. The information can thus be shared by anyone else that a user designates. In an office environment, for example, an administrative staff person can maintain an office address book that allows everyone to pick up the phone and request a specified person without having to look up the correct number for that person. The Web interface makes maintenance easy and practical.

A result of this approach will be less reliance on an actual physical unit. With integrated voice command the network is the repository for personal data, not the device itself. Thus, users will not have to program a device. The average cell-phone user now upgrades phones about every 12 months. If all pertinent information is maintained on the Web rather than in the phone, the user is spared many of the problems associated with changing.

FIGURE 1

Integrated Voice Control

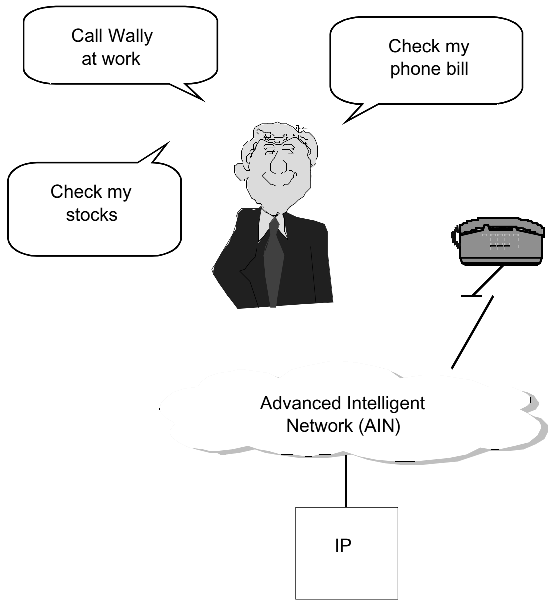

FIGURE 2

Meet-Me Calling/Conferencing

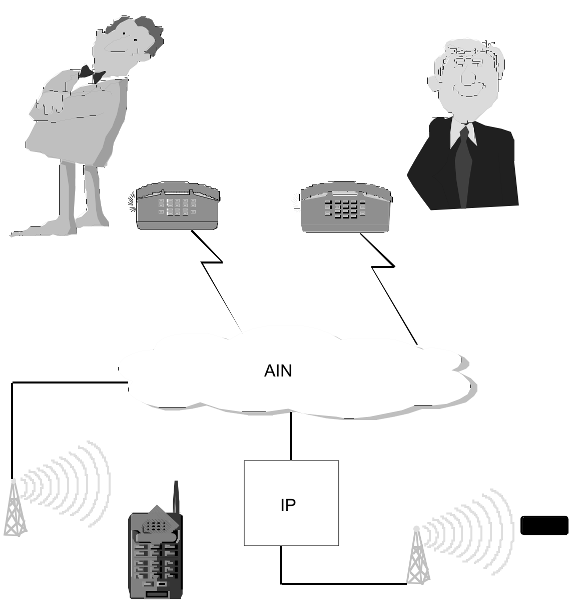

Voice-activated services bring value to the customer, which in turn increases customer loyalty. Increasing customer loyalty decreases churn, an expensive problem for telephony providers. In addition to saving money by keeping customers, voice-activated services facilitate usage of premium services and many more minutes of service, two ways for providers to generate additional revenues.

Meet-Me Calling/Conferencing

Another type of value-added service has been developed in prototype but has not yet been implemented. This service enables users to meet in the network for conferencing. Businesspeople who have been involved in a conference call and found that they needed to bring in someone else, but were unable to do so, will appreciate this service.

Meet-me calling allows two or more people to meet in the network (see *Figure 2*). The service is initiated with a series of special dual-tone multifrequency (DTMF) digits that page the parties or with a Web transaction that e-mails the people. Either way, the required parties receive a message requesting that they join a conference. The message provides a dial-in number and a personal identification number (PIN) for access. The application can dynamically add or drop parties as needed instead of maintaining a specified number of ports as required with current conferencing methods.

An example will show how the service works. Say that Randy receives a page requesting his presence at a conference. He calls the number given and enters the password. He can then communicate with Mark, who paged him, using Internet protocol (IP). After speaking with Mark, Randy may decide he needs some help from Oliver to solve Mark's problem. He enters a special series of DTMF digits, which places Mark on hold, and sees a menu that allows him to page Oliver.

Operator Redirect

The first two services outlined were revenue-generating services. The third, on the other hand, is a cost-avoidance service. With all the deregulation in the industry, the operator that a subscriber reaches is not necessarily the operator who can answer that person's questions. When people pick up the phone and dial out, in most cases they are connected to their local service company's operator. That operator has to spend time to understand what the caller really wants and to redirect the person to the proper place. This is a time-consuming and expensive process, and saving operator minutes can amount to significant cost savings. Thus, a service that can prevent such calls is of particular interest to the old regional Bell operating companies (RBOC).

Faced with this problem, one customer implemented an application using the multiservice platform that intercepts calls to the operator (calls dialed "0"), prompts the caller for a specific call request, then routes the call to the local operator if that is the proper person to answer the request (see *Figure 3*). Eliminating operator time to route requests for information has generated significant cost savings.

Conclusions

Many applications are currently being developed and implemented, and the three outlined here illustrate how they can generate revenue or save money while keeping subscribers committed to the service. In these ways, such value-added applications will keep telephony providers competitive in a very competitive industry.

FIGURE 3

Operator Redirect

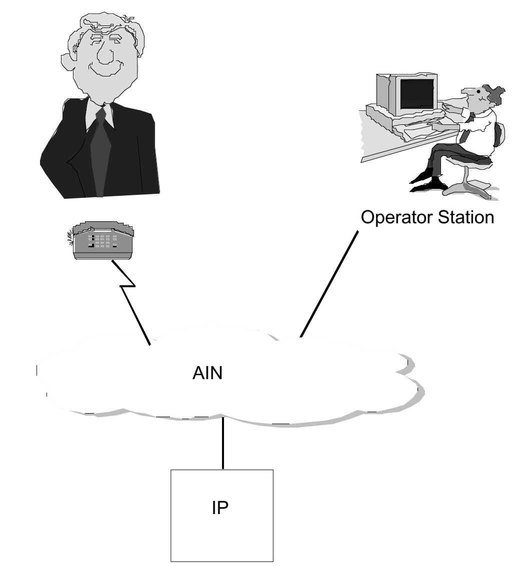

Optical Ethernet: Stirring the Imagination

Stephen Garcia
Director, Optical Ethernet Marketing
Nortel Networks

Eric McDonald
Senior Manager, Optical Ethernet Marketing
Nortel Networks

What if all corporations could connect buildings thousands of miles apart as easily as if they were just connecting two floors of the same building? What if every access and metro network was as simple, fast, and reliable as a local-area network (LAN)? What if metropolitan-area networks (MANs) and wide-area networks (WANs), with their multiple handoffs and protocol translations, physical distance, and hundreds to thousands of users, were as simple to manage as a single LAN today? Can you imagine the possibilities? Wouldn't it be revolutionary? The revolution is spelled OPTICAL ETHERNET.

Today's Corporate Networking Objectives and Challenges

According to a recent survey by Forrester Research[1], sixty-eight percent of Global 2500 corporations are looking for their information technology (IT) networks to provide competitive advantage. IT networks can provide corporations with competitive advantages in a variety of ways, including faster delivery of information, better access to customers and suppliers, support for next-generation applications, enhanced employee collaboration and productivity, and improved resource utilization

Corporations, however, face a variety of networking challenges when pursuing these objectives. Network complexity is increasing daily with the addition of new users, new nodes, and new network links. As traffic traverses the network from source to destination, it often undergoes a series of protocol translations, each of which adds complexity to the network. Network complexity represents the number-one networking problem faced by the Global 2500 today, as indicated in a recent report from Forrester Research[1].

The scarcity and expense of access bandwidth in today's networks represents an additional challenge for corporations (see *Figure 1*). Today, corporate LANs run at hundreds of megabits per second (Mbps). Similarly, WANs, built on optical systems capable of terabit speeds, are blazingly fast. Unfortunately, connections between the LAN and the WAN are extremely limited and have created the equivalent of an access bandwidth bottleneck. The average corporate site today uses a fractional T1 (64 kilobits per second [kbps] channels), a full rate T1 (1.5 Mbps) or, at best, a T3 (44.8 Mbps) to connect their LAN to the WAN. While 44.8 Mbps may sound like a lot at first, consider the case of a corporate network with 100 desktops. Assuming 80 of the 100 desktops use a full-duplex 100 Mbps Ethernet network interface card (NIC) and that 75% of the traffic generated by those 80 desktops is destined for the WAN, 12 Gbps can converge on the business' access link at the same time. While this represents an extreme example, even 10% of the traffic, 1.2 Gbps, represents far more than a T3 can provide. Compounding the problem of the access bandwidth bottleneck are the long provisioning times required for turning up new links—typically measured in months—and the inflexible nature of the service. Corporations are unable to scale their bandwidth linearly, and instead must scale their bandwidth in increments of T1s or T3s.

Network complexity and the access bandwidth bottleneck lead to a third challenge: poor network performance. Every time a packet undergoes a protocol conversion or requires a Layer-3 look-up, latency, jitter, and unpredictability are introduced into the network. Similarly, the access bandwidth bottleneck creates the need to throttle back the number of packets-per-second and to establish queues that add additional latency and jitter and degrade network performance. For existing applications and current networking paradigms, this level of performance may be acceptable. For chief information officers (CIOs) implementing new applications or trying to consolidate network resources in the WAN, today's network performance represents a critical obstacle.

A fourth challenge confronting corporations is the issue of IT staffing. As networks have become increasingly complex, the skills required to run the network have become more and more specialized. Whereas, in the past, network generalists could perform the majority of network jobs, today corporations must recruit, hire, train, and retain network specialists with detailed knowledge of specific network functions, such as security or the operations and maintenance of each specific technology used in the network—e.g., Internet protocol (IP), asynchronous transfer mode (ATM), and frame relay (FR).

FIGURE 1

Bandwidth Bottleneck

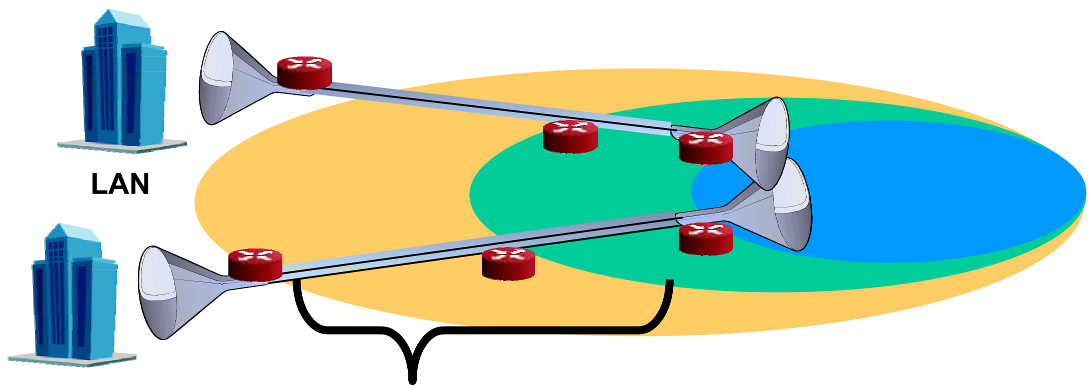

Optical Ethernet: Meeting the Corporate Challenge

Optical Ethernet allows corporations to overcome their networking challenges and utilize their networks to gain the competitive advantage that they desire. Fundamentally, optical Ethernet represents the combination and extension of two existing technologies, Ethernet and optics. This marriage takes the best of both Ethernet and optics and magnifies their capabilities to create a paradigm shift in networking and service models for both corporations and service providers. Whether offered as a managed service by a service provider or operated as a private network by the corporation, optical Ethernet transforms the corporate network into a key competitive advantage.

Today's Ethernet

The first personal computer (PC) with an Ethernet LAN adapter was shipped in September of 1982. Ethernet originally ran over thick coax and provided users with a shared 10 Mbps bandwidth connection. Ethernet soon progressed to running over unshielded twisted pair and offering dedicated 10 Mbps connections using switches. Today, Ethernet enables dedicated 100 Mbps to the desktop with 1 Gbps trunks, and, within a few years, industry experts predict 1 Gpbs to the desktop and 10 Gbps trunks. In the nearly 20 years since the first Ethernet shipment, Ethernet has become a widely standardized plug-and-play technology that is used in more than 90 percent of corporate LANs. Clearly, Ethernet has come a long way.

Today's Optics

Not to be outdone, optical technologies have come as far as Ethernet in arguably a shorter amount of time. It has been said that the advances in optical componetry are outstripping Moore's Law by eight times! Optical transmission speeds have grown from tens of megabits per second to 40 gigabits per second, and vendors have recently demonstrated the ability to put 1.6 terabits per second on a single optical fiber, using dense wavelength division multiplexing (DWDM). Optics' tremendous capacity as well advances in all-optical, or photonic, networking that eliminate the need for electrical regeneration have secured optics as the preeminent, lowest-cost-per-bit transport technology.

More than the Sum of Its Parts

Optical Ethernet, however, is more than the just Ethernet plus optics. Industry standards bodies as well as vendors have been developing optical Ethernet solutions that are higher in performance, more reliable, and more secure than merely running Ethernet over optics. Several standard bodies and industry alliances are working, for example, to create specific policing, traffic management, quality of service (QoS), and security mechanisms for optical Ethernet.

As a Layer-2 connectionless technology, optical Ethernet removes the addressing and other network complexity issues seen with IP and ATM–based networks (see *Table 1*). In addition, because it spans the LAN, MAN, and WAN, optical Ethernet removes the need for multiple protocol conversions that create significant management headaches for network operators. Furthermore, this simplicity extends into the provisioning and re-configurations of the network. No longer will businesses have to wait weeks or months for an additional T1 or spend time reconfiguring the network every time a change is made.

In addition to simplicity, speed is a key attribute of optical Ethernet. With the deployment of optical Ethernet, the bandwidth bottleneck is broken. Optical connectivity allows for access speeds of 10 Gbps—orders of magnitude faster than today's T3s. Bandwidth is also available in more granular slices. No longer are IT staffs made to jump from a T1 to a T3, when all they really need is another megabit of bandwidth. Optical Ethernet access links can be increased in 1 Mbps increments to deliver bandwidth from 1 Mbps to 10 Gbps or anywhere in between.

Optical Ethernet also greatly improves network performance. Because optical Ethernet is Ethernet end to end, it removes the need for multiple protocol conversions that introduce latency and jitter into the network. Lab tests conducted by Nortel Networks reveal that optical Ethernet reduces latency by more than 30 percent and reduces jitter by more than 90 percent when compared to traditional routed networks.

TABLE 1

Key Network Attribute	Today's Networks	Optical Ethernet
Complexity	High: • Layer 3 • Multiple protocols from LAN to WAN • Fixed design	Low: • Layer 2 • Ethernet from LAN to WAN • Flexible design
Access Bandwidth	• Fractional T1s, T1s, T3s • Weeks to months to provision • Don't scale linearly	• Up to 10 Gbps • Hours to days to provision • Linear scalability
Performance	• Unpredictability • Latency • Jitter	• LAN performance end to end
Staffing	• Network specialists for each protocol	• Network generalists

Moreover, optical Ethernet topologies enable much greater reliability than today's access networks can provide. For example, an Ethernet over resilient packet rings (RPR) solution provides less than 50 milliseconds of delay in event of a catastrophic failure, such as a line cut. This high availability guarantees up time in networks that deliver mission-critical applications. The reduced latency and jitter in an optical Ethernet network also delivers the performance necessary to run time-sensitive, multimedia applications such as videoconferencing.

Finally, optical Ethernet is significantly less expensive than today's networks. The cost of an Ethernet interface card, for example, is a fraction of the cost for an ATM or packet over SONET (POS) interface card. In addition, the optical Ethernet's Layer-2 simplicity allows network operators to eliminate multiple private virtual circuits and reduce staffing levels. Nortel Networks, which is converting its own internal network—a network that spans six continents and supports more than 50 million voice minutes and 1.3 petabytes of data traffic per month—estimates that its cost savings from optical Ethernet will be between 30 and 40 percent.

Why Now?

The convergence of multiple factors has created the market discontinuity that optical Ethernet is now exploiting. First, as was mentioned, corporations are continuing to choke on the increasing complexity of their IT demands. Second, competitive price pressure is forcing service providers to move up the value chain away from undifferentiated products and services and toward value-added services. Third, several critical optical Ethernet technologies have matured to the point that optical Ethernet is now ready for commercial adoption—Ethernet has become carrier-grade, and the availability of fiber in access networks has increased[2]. Fourth, the Internet has become a trusted business tool.

Together, these factors are driving significant demand for optical Ethernet. The Yankee Group recently forecasted a cumulative $5 billion market and a compound annual growth rate of 90 percent for optical Ethernet equipment between 2001 and 2004 (see *Figure 2*).

Building an Optical Ethernet Network

There are three basic building blocks that can be used to construct an optical Ethernet network, whether private (corporate) or public (service provider). Each of these building blocks delivers its own set of advantages. Specific optical Ethernet network topologies will vary depending upon the needs of the network operator and in many cases may include more than one type of building block (see *Figure 3*).

Ethernet over Fiber

The first building block is Ethernet over fiber. This implementation is simple electrical-to-optical translation over non–SONET–based fiber. Only Ethernet services are available with this type of implementation (no time division multiplexing [TDM] services), and these services are somewhat limited to campus-type installations because of a less than 70 km distance restriction but can also be used in hybrid optical Ethernet solutions. This implementation is a very low-cost alternative because all interfaces are Ethernet only.

Ethernet over Resilient Packet Ring

A second optical Ethernet building block is Ethernet over RPR. While RPR standards are still being developed, many customer specific implementations are carrying traffic today. Ethernet over RPR requires the availability of SONET infrastructure but solves the traditional problem of SONET bandwidth "waste" because in Ethernet over RPR the SONET ring's bandwidth is utilized in both directions instead of only one. Ethernet over RPR carves up a portion of the SONET ring to create a multiservice ring allowing for pure Ethernet services as well as TDM, ATM, and IP traffic. Ethernet over RPR takes advantage of SONET's link-layer protection and provides less than 50 millisecond fail-over. The RPR can span thousands of kilometers and hence provides a flexible optical Ethernet solution for any number of implementation scenarios.

Ethernet over DWDM

Ethernet over DWDM utilizes DWDM as its core transport versus SONET. It delivers massive bandwidth. This implementation is by far and away the lowest cost per megabyte solution and is most appropriate where the application(s) dictate very high bandwidth requirements (such as storage-

FIGURE 2

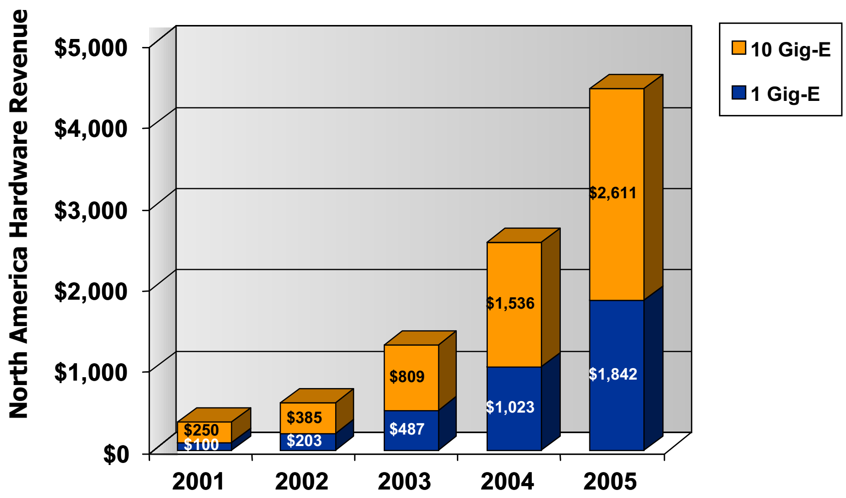

Source: Yankee Group, 2001

FIGURE 3

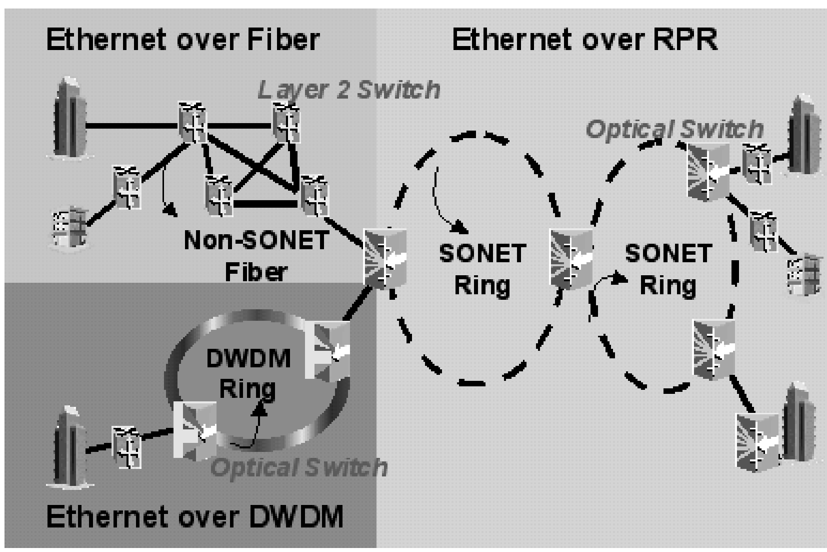

area networks [SANs] and data centers) or in areas where SONET is not available or where fiber availability is limited and must be most effectively used. Ethernet over DWDM is bite-rate and protocol independent, which allows it to work in almost any network scenario.

Benefits to Corporations and Service Providers

Corporations and service providers will derive a number of tangible benefits from the adoption of optical Ethernet. Corporations that adopt optical Ethernet will immediately benefit from improved network performance (more bandwidth, reduced latency and jitter, greater network predictability, and network utilization) and reduced capital and operational costs. Moreover, because optical Ethernet creates LAN–like performance in MANs and WANs, corporations that deploy optical Ethernet will be able to consolidate network infrastructure and applications, thereby realizing significant cost savings. Let us assume for example that Corporation XYZ has two primary sites, one in Boston and one in New York. In all likelihood, Corporation XYZ is maintaining a separate data center and a separate Internet service provider (ISP) connection for each corporate site. By deploying optical Ethernet, Corporation XYZ can consolidate both data centers and ISP connections at one site, e.g., Boston. Employees in New York will still have access to all the same applications, at the same levels of performance, by connecting over the corporations' optical Ethernet network. The corporation, however, will save significantly by eliminating the cost of a data center and ISP connection.

For service providers, optical Ethernet represents the opportunity to continue the evolution of their services up the value chain. At the end of 1998, service providers' wholesale bandwidth prices were falling by 50% every 18 months. Since the middle of 2000, prices have been dropping by 50% every six months.[3] Not surprisingly, service providers are looking for new services to increase their revenues, expand their customer base, and deliver competitive differentiation in their service offerings. Optical Ethernet enables these capabilities.

Optical Ethernet services generally fall into two categories, connectivity and value-added services. Connectivity services look very much like services that service providers are offering today—services such as private line, aggregation and backhaul, and LAN extension. While connectivity services may cannibalize service providers' existing revenues in the short term, optical Ethernet's decreased operating expenses will help to offset this effect and maintain service providers' margins on these services. Moreover, optical Ethernet connectivity services lay the foundation for service providers to offer differentiated value-added services.

The market for value-added services is large and growing (see *Figure 4*). Not surprisingly, service providers are racing to offer these services. Unfortunately, today's networks are capable of only the most basic value-added services, such as Web hosting. Optical Ethernet, however, provides the network-performance improvements necessary to make value-added services a reality for service providers. Decreased latency and jitter, and better network predictability makes offering real-time and mission-critical services such as voice over IP (VoIP), storage-area networking, and disaster recovery true possibilities for service providers. Furthermore, service providers will be able to differentiate their services with better service-level agreements (SLAs). Tests conducted by Nortel Networks show that optical Ethernet improves SLA response times over ATM by 44 percent.

Summary

Optical Ethernet delivers today what could only be imagined yesterday. By combining the ubiquity and reliability of Ethernet with the speed and economics of optics, optical Ethernet creates a networking paradigm shift. Optical Ethernet solves the critical challenges faced today by corporate IT managers and service providers. Optical Ethernet's simplicity is found in its easy network management and provisioning. Additionally, optical Ethernet removes the complexity of multiple protocol translations that can add latency and jitter into the network. Optical Ethernet is not only fast; delivering speeds from 1 Mbps to 10 Gbps, but it also provides for linear, predictable bandwidth management. It removes the limitations seen in many of today's access and metro networks that can bottleneck bandwidth availability in the network. Finally, optical Ethernet delivers mission-critical reliability in fail-over times, physical security of fiber, and enabled reliability in disaster and recovery and network-area storage applications. The coming optical Ethernet revolution will be limited only by our imaginations, eclipsing the changes that Ethernet or optics ushered in and delivering, in one solution, a fast, simple, and reliable network.

Notes

1. Forrester Research, *The New Enterprise Network*, January 2001.
2. According Dain Rauscher Wessels, 35 to 45% of all enterprises are within 1 km of fiber.
3. *Red Herring*, January 16, 2001.

FIGURE 4

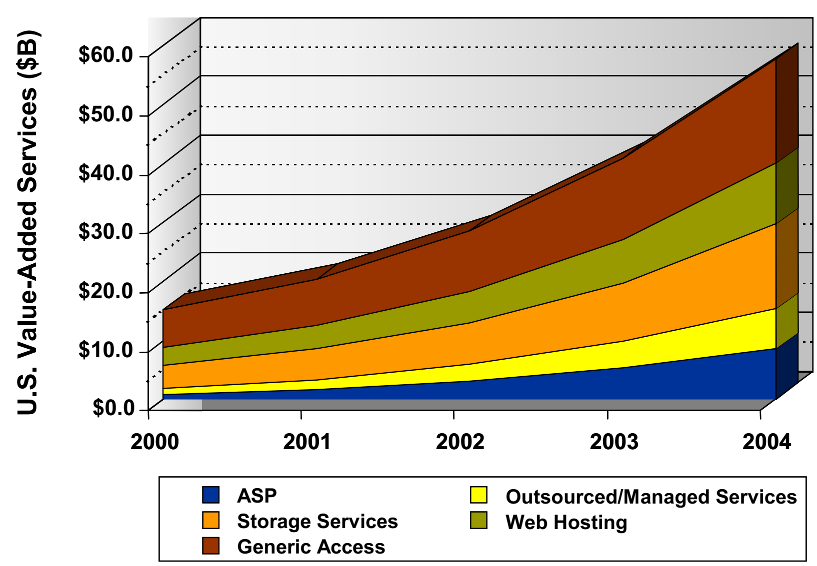

Source IDC and Probe Research

Understanding the Design Challenges for Residential Gateway Chips

Bernard Goffart
Product-Line Manager, Gateway Products
Alcatel Microelectronics

Abstract

Key customer-premises equipment (CPE) gateway chip design issues include technology choices, interoperability, and cost reduction. Designers have choices in determining which local-area network (LAN) and digital subscriber line (DSL) access technology to use and how to integrate the needed functions at the lowest possible cost. This paper will cover how chip designers can enable interoperability between different mediums from residential gateways to peripherals; how voice, data, and real-time video digital content can be transported reliably and securely; and how home LANs can use existing cabling plus ease of set-up, configuration, and management. Finally, the reader will gain an understanding of how the design of gateway chips enables upgrading to the integrated access device (IAD), allowing the delivery of future value-added products from service providers. As the IAD become increasingly sophisticated, the addition of new functions and additional capabilities will not only fuel the installation of home gateways, it will also stimulate the provision of many new services to the home user.

Introduction and Overview

Networked homes of the future are becoming closer to reality in large part due to breakthroughs in gateway technology. The day when appliances, security systems, telephones, computers, and entertainment systems are connected will soon be here. As rapidly increasing numbers of households with multiple computers and Internet access demand higher levels of telecommunications speed and data capacity, DSL connectivity is entering the residential space. At the same time, in response to market demand, satellite and television cable operators have launched data over cable service interface specifications (DOCSIS) digital communications systems to offer fast Internet access. Access technology previously reserved for business applications is streaming into the home.

Connecting the home of the future presents challenges unlike those found in the business environment. The average household will not have a staff of systems professionals available for maintaining a network and resolving complex problems. Legacy equipment will need to be refined for ease of installation, configuration, use, and maintenance as well as affordability. Furthermore, the diversity of applications and systems found in the home makes demands on a network unlike those of a business. Residential gateway products capable of supporting residential requirements have started to enter the market, and many more are expected to enter the market over the next several years.

Example Applications

The Telecommunications Industry Association (TIA) has identified initial residential gateway applications to include the following:

- Telecommuting
- Internet access
- Distance learning
- Telemedicine
- Video telephony
- Home appliance management and integration
- Security systems management
- In-home power regulation and management
- Automated meter reading
- Neighborhood cordless roam phones
- Video delivery and distribution
- Virtual videocassette recorder (VCR) and video on demand (VOD)
- Video intercom
- Compact-disc (CD) jukebox
- On-line advertising and electronic catalogs

Given that the home market is known for being price sensitive, the first-out products will likely not be the highest bandwidth-consuming products that could be connected to the LAN for the aforementioned applications. As semiconductor technologies continue to provide more integration opportunities and improved circuit performance at lower costs, these products will become available and be joined by even more practical applications. The design engineer is therefore best advised to consider application flexibility with the design so that additional product functionality and value may be added over time.

Gateway Design Characteristics

Future gateways will not depend upon personal computers (PCs) and will be able to offer the 99.999% availability that telecom service demands and PCs cannot provide. The mainstream residential gateway will not be dependent upon a PC. Rather, it will be a standalone box able to route information at high speeds and provide hub services with always-on Internet access. Therefore, it requires Internet isolation and protection functions to avoid hacking. The residential gateway will be designed to be durable, with a 10-year life span, low failure in time (FIT), and recovery procedures in place. It will need to support a remote management capability. It will support multiple services starting with telephony and VOD. These gateways will be portals to a number of new applications based on the high-speed on-line connectivity.

Example of Expected Broadband Piping Improvements

We use asynchronous digital subscriber line (ADSL) as the example (see *Figure 1*). The most exciting feature of ADSL is that it provides an "always available" data channel to the home or office. And with asynchronous transfer mode (ATM), it is designed to handle multiple channels with varying quality of service (QoS). ADSL allows for the development of services that were previously uneconomical. Push as well as pull applications are becoming available. The main interface box between the broadband pipe and home networks will be the residential gateway or IAD when multiple services are enabled.

The ATM–based multimedia pipeline is positioned to provide expanded services. There is great opportunity to deploy exciting new applications, but there is room for ADSL broadband pipe design improvements to support expanded services and customer satisfaction.

A few examples will be given. DSL deployment and troubleshooting must be improved. For the DSL service provider, this includes automated logistics for deployment, line, and service qualification. Additionally, the user is provided remote monitoring and diagnostics capabilities for cases where there is difficulty during installation or during later operation. This supports user self-install and eliminates the need for on-premises presence for installation or maintenance support. Remote CPE and central office (CO) operations, administration, and maintenance (OAM) and diagnostic tools will be applied. Eliminating the plain old telephone service (POTS) line splitter helps to simplify self-installation, avoiding the need for additional equipment to be installed at the customer premises. Unfortunately, performance quality of the POTS line is achieved only by inserting a micro filter in front of each phone. With anticipated future deployment, an all-digital loop approach will help to reduce the deployment costs. In this case, the POTS line is replaced with ADSL service, providing for all-digital transportation of the voice conversation.

Forthcoming improvements are expected for the ATM/ADSL multimedia pipeline infrastructure in radio frequency (RFI) robustness. To improve the operation of ADSL with HomePNA, and where multiple lines in the same bundle are used, enhancements will include better-controlled impedance and spectrum management. Residential gateway designers will welcome such improvements in the ATM/ADSL multimedia pipeline infrastructure.

Vision of the Evolution of Gateway Functions

IADs will need to evolve substantially over time to meet the gateway needs of residential users. Today's IAD consists of a PC network (for example, PC–to–PC) that includes shared Internet access, a printer, and possibly a firewall to the public Internet. Access to the Internet may be via an analog dial-up modem or a broadband modem. As reported by Kinetic Strategies, more than nine million homes in the United States have broadband access at this time. Early adopters are using voice over Internet Protocol (VoIP) Web phones and streaming audio for music, infomercials, and video for train-

FIGURE 1

ADSL Multimedia Pipeline

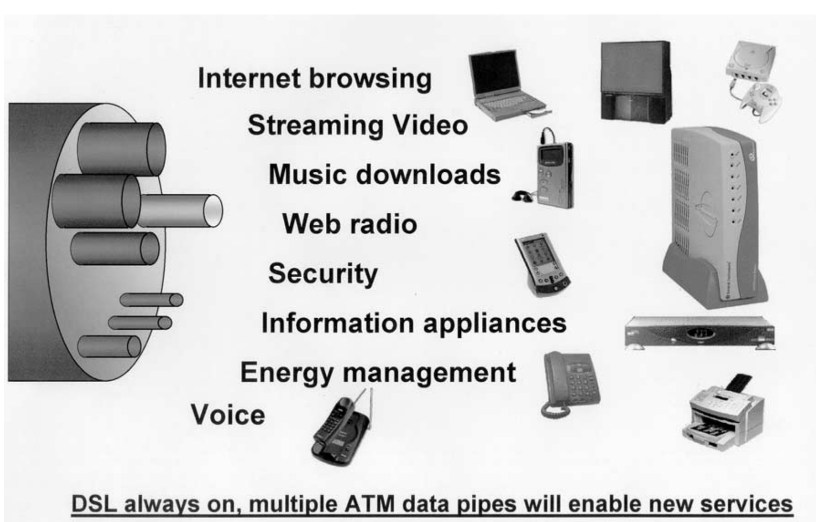

ing programs. There has been little demand for a residential gateway and LAN.

In phase two, however, the basic PC IAD will begin to transition to a dedicated IAD capable of processing integrated data and voice as well as managing a home network. Early adopter residential users will begin to access high-definition television (HDTV) broadcasts via the Internet, watch high-quality videos on-line, and have brown goods such as the television, VCR, and stereo equipment connected. Phase two will have a mix of residential users—some still using the PC as the IAD, but a growing segment of users becoming dependent upon a dedicated IAD and residential LAN.

Phase three is predicted to start within about five years from this writing and will see the rapid ramp-up of fully integrated services that provide universal connectivity, with a port available on all types of devices. Home security will be integrated and white goods (appliances) and brown goods (e.g., entertainment) will be connected, and mobility integrators will offer new products. The delivery of services from vertical suppliers—such as monitoring services, security services, and financial applications—are expected to have wide acceptance. Applications that extend the office workplace into the home will likely be the initial applications that are sufficiently compelling to make the residential gateway and LAN investments. Home control systems, security, medical support, and home entertainment distribution follow.

During phase three, providers of residential gateway products will find users requiring a range of performance and capability depending upon the application. At the high end of residential gateway hardware, the small office/home office (SOHO) will include a multi-host DSL router with eight to 24 lines of voice over DSL (VoDSL). This model will allow a number of users to be attached to a single server. The strictly residential model will provide VoDSL via a low-end router with two to four lines. For those households wanting only data services, a DSL router capable of connecting more than one computer directly to a DSL line, with integrated management and security features, will provide a moderate level of features. Future IADs will also include VoIP services at this level. The single host PC platform will no longer need a gateway, with a DSL modem operating from the internal peripheral component interconnect (PCI) bus or bridged externally to a universal serial bus (USB) or Ethernet.

Phase four will be the maturation phase similar to how we now consider PCs—not an every-home item but quite pervasive. Existing homes will have been retrofit with the preferred LAN technologies, and nearly all SOHOs and new homes will be built with the residential LAN. Users will have many choices of services and activities accessible via the residential gateway. They will be able to choose from many microprocessor-operated home appliances plug-and-play ready to be networked via standard interface ports.

Introduction to Design Challenges

Complexities of the interworking of wide-area networks (WANs) and LANs make the design of products for residential applications an interesting challenge. Exchange of information must be processed between two different worlds, the wide-area access network and the local network. On the access side, a number of protocols must be addressed. Today, telecommunications companies, mainly incumbent local-exchange carriers (ILECs), are ATM–centric. They and Internet service providers (ISPs) are adopting Internet protocol (IP)–based solutions. Since these technologies are used in different segments of the data path architectures, technologies for addressing both IP and ATM framing characteristics are needed. These technologies are known as encapsulation techniques. Thanks to several interoperability initiatives, the stabilization process is currently underway. Access authentication and administration is also done with varying technologies. Most deployed methods are using the legacy point-to-point protocol (PPP) developed for dial-up access. Methods are adapted to Ethernet (PPPoE) and ATM (PPPoA). The addition of new voice-over-packet services, such as VoDSL or VoIP, is also adding new dimensions to the interworking of networks. Within two to five years of this writing, it is expected that new services will be added that will extend the IAD with voice, data, and video capabilities. The impact of the introduction of IP version 6 (IPv6) into the backbone network should be taken into account. Mobile devices are expected to be the first to make use of IPv6 to the terminal. The evolution of Ethernet to the last mile will significantly impact the residential gateway architecture.

Residential Gateway Design Issues and Solutions

Connectivity support for residential gateways requires working in two worlds, ATM and Ethernet. The transportation mechanism is ATM while the payload is Ethernet requiring interoperability between the two. There are also multiple alternatives for session control—namely PPP and virtual private network (VPN). The flavors for broadband (ATM and Ethernet) can be compared.

PPPoA (PPP over ATM) uses one virtual circuit (VC) per session on an ATM adaptive layer (AAL)–5. Powerful ATM QoS tools may be leveraged. Point-to-point tunneling protocol (PPTP) for remote session set-up is required. Multiple sessions require support for multiple VCs.

For PPPoE (PPP over Ethernet) one VC for an unlimited session number is used. There is RFC1483 encapsulation over ATM, and only a best-effort QoS is supported. ATM is used as a transport mechanism between the CPE and DSL access multiplexer (DSLAM).

Guidelines for access service provisioning include an allowance for a simple set-up process and provisions for authentication, authorization, and accounting (AAA). Compatibility should be maintained with a layered approach, depending upon the country, with regulatory plus network service provider (NSP) mapping on installed layer two. We also recommend bandwidth shaping with differentiated capability for QoS and that proper security features be provided. The design should be based upon ubiquitous access technology as close as possible to V90 dial up or broadband.

Quality of Service

The control of the QoS over the entire chain access, gateway, and LAN will promote the emergence of new applications sensitive to latency and throughput (audio, video, conferencing). It is known that compression and echo-cancellation audio algorithms, and the data packetization process, require intensive signal-processing support. Particular attention must be paid within the CPE to minimize the impact on the overall delay. Jitter must also be controlled. For example, telephony over packets (VoIP, VoDSL) is transported over a dedicated bus general control interface (GCI) or a combination of pulse code modulation (PCM)/serial port interface (SPI). The signal is passed over a classical pair of subscriber line interface circuits (SLICs) and audio-processing circuit devices that implement the digital signal processor (DSP) functions and allow for the interface to the legacy analog telephony equipment. Home LAN support of QoS will eliminate the need for POTS interfacing. As an illustration of this concept, telephony can be transported over a LAN network with specific LAN phones placed on the line. These phones will interface directly with a set of loudspeakers and microphones connected to a LAN medium after digitization, coder/decoder (CODEC) echo cancellation, and packetization; or they interface through an RF link to a cordless phone. Furthermore, the all-digital loop path DSL gateway may eliminate the need for the dual POTS/DSLAM equipment at the CO side by relying on the voice-over-packet transport to the LAN phone.

All-digital loop refers to the access technology where all information, including the telephony signals, is transmitted as digital data up to the customer. For ADSL or VoDSL, this implies that, in contrast to the traditional "overlay scenario" where POTS or ISDN are transported "out of band," the telephony signals are transported in-band. The telephone signals are digitized on the CPE such as on the LAN telephone and forwarded to the access node through the broadband access network.

Network Isolation and Security

Network isolation and security encompass IP network access translation (NAT) or port access translation (PAT), encryption, and a firewall kit. Support for specific applications is mandatory. Typical solutions are providing application language gateways (ALGs) that support specificities for approximately 30 to 35 applications. To understand the complexity of a NAT/PAT, one needs to understand that for many applications, local IP addresses are included with the payload of IP packets. Examples are H.323 and Netmeeting. To make proper NAT, the tools must be capable of making address substitution within the payload. Other functions included in the firewall kit are a standards-based packet-filtering engine, proxy server, socks server, stateful multilayer inspection (SMLI), and compliance with International Customer Service Association (ICSA) firewall certification. For CPE applications, the code size must be managed.

Other considerations for firewall engine security are access policies/access restriction and protocol disabling. Remote login user configuration should be supported.

Local-Area Networks

There is a plethora of wire and wireless LAN implementation technologies in consideration for becoming the preferred residential LAN. Some of the technologies, such as Ethernet and Firewire, listed in *Table 1* are broadly applied, while many others are in an early stage of development. The table shows advantages and disadvantages that can be expected from the technologies if used for residential LAN applications. Some of the quantitative information is somewhat fluid in that new versions of standards are in development. For example, a new version of the standard for HomePNA is in development for 32 Mbps, as well as a new version for IEEE 1394 for 1,600 Mbps and a 1,000-meter range.

LAN Medium

The distribution and communications medium will likely include wire and wireless. The choices for use of existing wire are the home phone line, supported by the Home Phone Network Alliance (HPNA) or power line supported by the HomePlug Alliance. Typically requiring new wiring would be Ethernet (IEEE 802.3) unshielded twisted pair (UTP) Cat 5 cable, Firewire (IEEE 1394), or structured cabling where multiple wire types are within a cable. Wireless standards viewed as having the highest chance of wide acceptance for residential application are HomeRF, IEEE 802.11b, and Bluetooth. With respect to the transmission speeds for wireless that are shown in *Table 1*, they are reduced as distance increases, but wireless provides the greatest flexibility for users.

Ideally, the overall guiding principles for the design of the connectivity medium infrastructure would be to use the existing cable as is. Deployment of new cabling for home networking is costly. The goal is also to achieve simple setup procedures and to support QoS tools such as opening the LAN to real-time or near–real-time applications.

LAN routing traditionally has been limited to business applications. The building of LANs is a complex task. Its maintenance is usually expensive in that it requires highly technical personnel. The services rendered in a business environment justify the investment. For a home application, the same level of cost is not affordable. A number of efforts are underway to improve the ease of configuration and maintenance.

The data-transmission capabilities of existing cables in homes have been investigated. Several methods are used for deployment. The industry has initiated standards. For example, the HPNA makes use of the phone network, HomePlug makes use of the power-line network, and wireless solutions—such as Bluetooth, IEEE 802.11, HomeRF, and HiperLan—are either already in use or anticipated for home application. LAN hardware must be designed for easy deployment, configuration, and maintenance of the home network.

The diversity of appliances destined to be hanging on the home LAN is increasing. LANs need to be able to support new types of applications and new types of services. Typically the "best-effort" philosophy of Ethernet protocol is a major limitation for streaming applications as well as for applications requiring latency control.

The availability of residential gateways capable of supporting future services will depend upon their ability to respond to the challenges described. The diversity of data media inside the home brings new challenges for system partitioning. Integrated micro switches, including media access controller (MAC) and physical (PHY), are available on the market for Ethernet unshielded twisted pair (UTP). The availability of these components justifies a system interfac-

TABLE 1

Alternatives for LAN Networks

Technology	*Performance*	*Advantages*	*Disadvantages*
Wired Medium			
HomePNA	1–10 Mbps	Compatible with POTS and ADSL	Insufficient number of phone jacks in typical home
HomePlug	10–14 Mbps	Typically sufficient outlets available	Resolution needed with some local regulations that forbid modulation in the required band.
IEEE 802.3 Ethernet	10–100 Mbps	Readily available	Most cases need new wiring; cost
IEEE 1394 Firewire	100–400 Mbps	Throughput needed for multimedia; isochronous	Limited to 5 meters
Wireless Medium			
HomeRF	1.6–11 Mbps	DECT–like (already shipping); access up to 127 devices	Limit of four simultaneous handsets; 2.4 GHz radio band is noisy
IEEE 802.11a Wireless LAN	24 Mbps	Throughput; up to 1,000-foot range	Cost; power consumption
IEEE 802.11b Wireless LAN	11 Mbps	Throughput; up to 1,000-foot range	2.4 GHz radio band is noisy
Bluetooth	1 Mbps	Cost; low-power transmission	Throughput; only 10-meter range; dial-up to use piconet; radio band is noisy
HiperLAN	20–54 Mbps	Throughput	Power consumption

ing at the media independent interface (MII) level. Another justification to interface residential gateways at this level is the evolving consensus among the industrial community to develop PHY devices for alternative home mediums at the MII level. For instance, there are such devices on the market for standards such as HPNA, 802.11, HiperLan, Bluetooth, and HomePlug. After a consolidation period, the range of media and standards will narrow, and residential gateways will embed the PHY part.

Protocols

The broadband pipeline infrastructure supporting a multiplicity of standard protocols for packet transmission is first described. This is followed with an example of software architecture appropriate for a residential gateway system.

There are various methods used for voice-over-packet technology (see *Figure 2*). These include using an ADSL PHY with VoIP or voice over ATM (VoATM). Some are still working on voice over T1 (VoT1) via AAL–1. Another method is channelized VoDSL (CVoDSL), and then the familiar lifeline bus over copper line. In the context of the example provided in this material, only the VoDSL and VoIP will be covered.

In this example architecture, voice and data information are entering from the DSL physical layer through the Utopia interface block, with the traffic being ATM cells. After segmentation and re-assembly two ATM channels are separated. The AAL–2 channel is used for voice (VoDSL) and the AAL–5 channel is used for data. For the AAL–5 traffic, after decapsulation of the ATM cells' content, the resulting Ethernet frames are sent to the client target using the MII interface. This interface transfers Ethernet frames generated by the MAC to the physical device that will send packets to the home network. One will note that VoIP traffic is captured within the transmission control protocol (TCP)/IP stackware and sent to the audio processing part. The VoDSL traffic, being separated at the ATM level, is directly sent to the audio-processing part.

Audio processing and call processing set up the telephone calls and maintain the session throughout the call. Depending on the type of technology in use (VoDSL or VoIP), different technologies and standards are in use to implement the required functions. Referring to the left side of *Figure 3*, signal processing–intensive functions are performed. These functions include adaptive echo cancellation (AEC), voice compression, signal mixing, activity detection, and others. The system services are depicted on the far left and include entities such as the management information base (MIB) and the CTRL–E to drive the ADSL chip.

To allow device customization and application integration, a number of application program interface (API) functions are provided. These API are providing access to various elements of the device architecture.

Gateway Application Support and Integration

A residential gateway product is not complete by just addressing the protocol requirements. Application drivers and software must be available for a complete solution (see *Figure 4*). There are a number of companies specializing in

FIGURE 2

Voice-over-Packet Technology

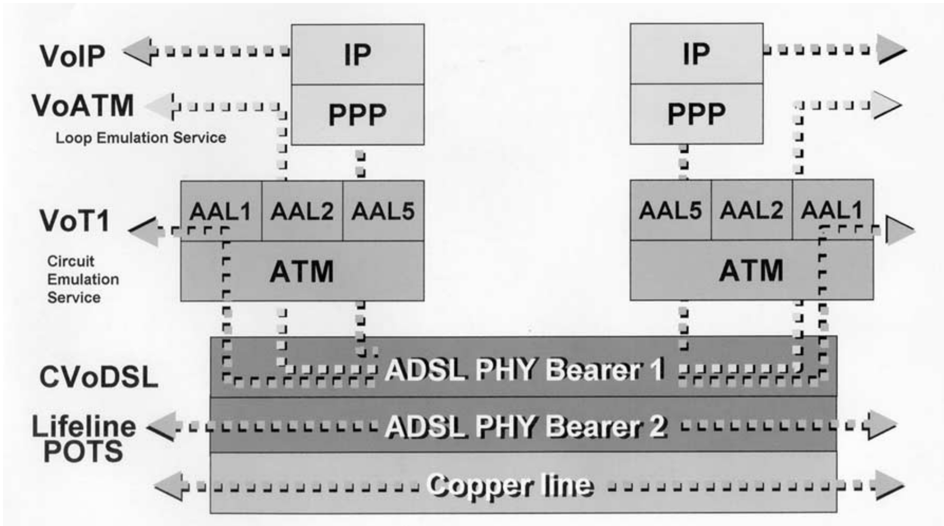

software product development for residential gateway applications, and such companies can provide these products. However, the chip and system developer must plan and test for interoperability of all components of the system.

An example block diagram for a design implementing home audio/video interoperabilty (HAVi) architecture is shown in *Figure 5*. The API is established for some of the applications with native binding and others for Java binding. A Java-based virtual machine provides a framework for the integration of service- or device-specific applications. While residential gateways may initially be called upon to service a very limited set of applications, standard interfaces promulgated by the open services gateway initiative (OSGI) and implemented in the system software architecture will allow for future applications to be added such as illustrated in *Figure 5*.

Design Implementation

At the board level, the goal should be to eliminate external components for cost reasons. Try to simplify board routing and use cheaper printed circuit board material with the reduction of board-level high-frequency signals. Consider using modules to simplify manufacturing.

Solutions at the chip level can help the design meet such goals with the integration of interfaces such as PCI/USB/

FIGURE 3

Software Voice plus Data Architecture

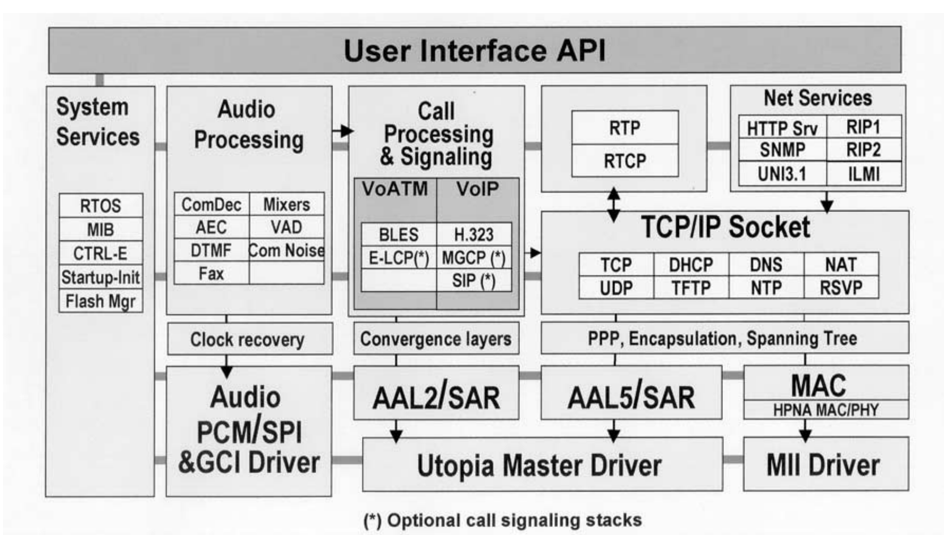

FIGURE 4

System Solution with Software

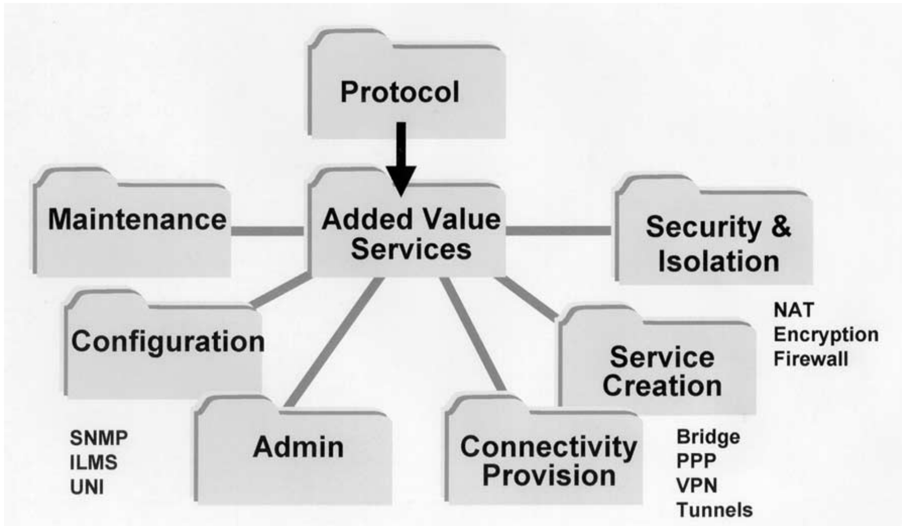

ATM/MII, integration of the access and LAN PHY, memory, and external filters. We see silicon integration occurring in four phases as illustrated in *Figure 6*.

The first phase is the exploratory phase, where the gateway is designed around the PCI bus with standard components. Application-specific integrated circuits (ASICs) have been used for glue logic and the PHY. While the standards are settling out, this phase would be used for the definition and check-out of interworking and software closure. This phase was to be followed by the availability of application-specific standard products (ASSPs) for the access modem and include the application of communications and network processors. During phase two, standards are reaching maturity, allowing developers to see stabilization of interworking protocols. Feature evolution is made possible by allowing for software plugs and API standardization.

Phase three is anticipated to be the mass production phase with complementary metal oxide semiconductor (CMOS) technology providing for further integration opportunities. Actually, the complete system will be integrated on a single chip. This system-on-chip (SoC) ASSP will include processor, access modem, and an embedded LAN PHY. In this phase, open architecture definition will expand application support from third-party software developers.

FIGURE 5

HAVi Device

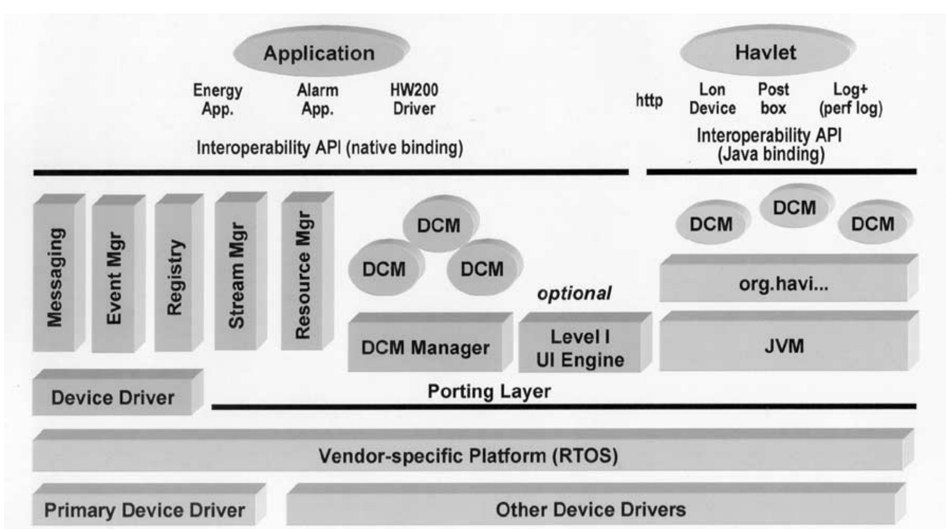

FIGURE 6

Hardware Architecture Evolution

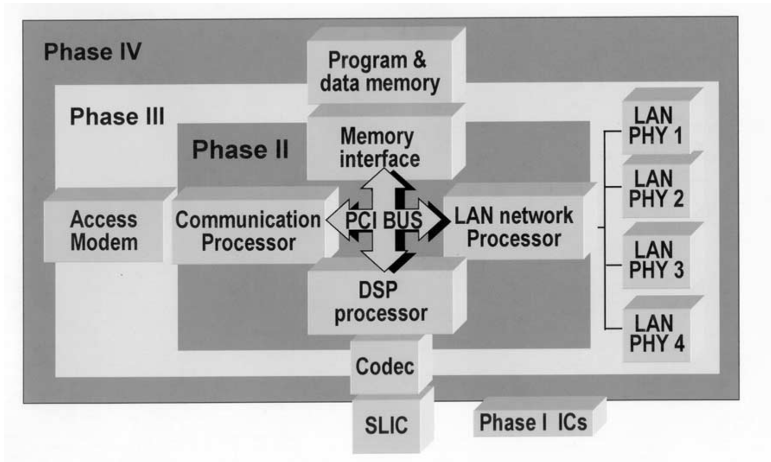

More cost reduction will be called for in phase four, with an integrated circuit (IC) in 0.13 or 0.11 CMOS technology that includes a modular architecture with communication ports, multiprocessor integration, phase-three external component integration such as glue logic, memory, and multi-PHY on-chip support. This design will be optimized over the phase-three design with low-cost software stacks.

A reliable semiconductor supplier will provide a complete development solution consisting of not just silicon, but also a reference design board with schematics, software drivers, evaluation software, documentation, full customer support, and a design example based upon the reference board.

Examples of a gateway design are shown in *Figure 7*.

Conclusion

To successfully address the large residential gateway market opportunity, we have concluded that it will be best to leverage legacy telecommunications access, evolving new LAN technology, and the potential of the Internet.

We have described the design issues and solutions for complex technological challenges of sophisticated interworking functions and multiple services in an environment of tough commercial targets. To achieve worldwide interoperability, industry-wide collaboration is mandatory. To facilitate system design with ever-changing standards and the advancement of product functionality, the silicon solutions must prove flexibility for system adaptation. A silicon-only solution is not a complete and flexible solution. A complete design kit solution is called for that includes much of the system software.

FIGURE 7

Wireless 802.11b plus Bluetooth Reference Designs

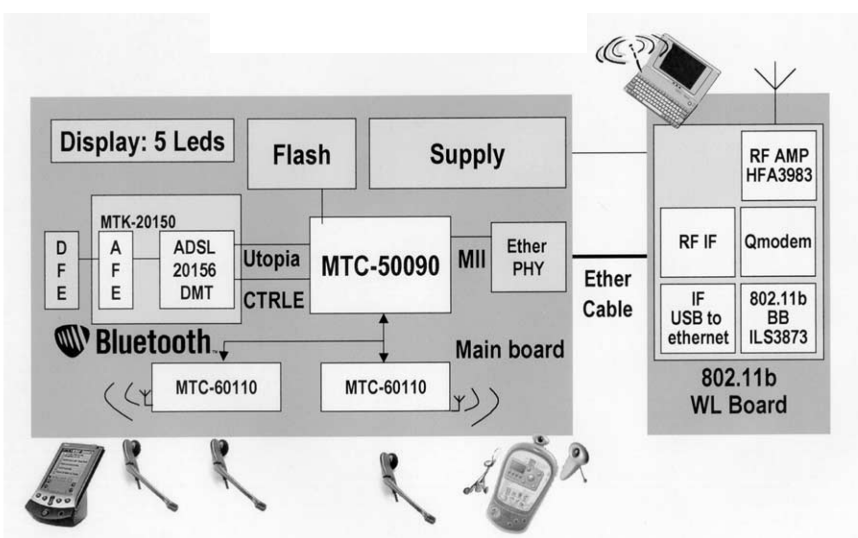

The Light Ahead: Integrating Optical Components on Silicon

Robert Green
Vice President, Business Development
Bookham Technology plc

Despite the recent slowdown in the global economy, the pressure for innovation in the communications industry has not abated. The demand for bandwidth is still growing, but the challenge no longer ends there. The competitive communications market of today is calling for minimized capital expenditure, but more importantly, it is also asking for increased flexibility in the networks on which it relies to generate revenue, and hence increased returns on its investment. Network providers want greater functionality; they need this to be able to deliver bandwidth at the right price, in the right place, and at the right time. These factors require that the infrastructure be used to its maximum potential. In the quest to satisfy the increasing demand for greater network capability, fiber-optic cable has been deployed aggressively and the push is now on to make maximum use of that investment, particularly in the current difficult climate.

The solution to these pressures is increasingly being found in increasingly integrated optical components, an area of the industry where the most noticeable advancements are emerging. This market has always been customer-led, but the nature of customers' demands is changing and creating greater competition. To meet the criteria for greater flexibility, the industry is leaning toward customized products with greater functionality. By integrating many optical functions as closely as possible, device manufacturers are beginning to offer themselves as one-stop solutions shops, not only producing components, but also a range of customizable optical subsystems.

Today, the optical components industry has solved the production and manufacturing problems of old by directly applying techniques from the semiconductor industry and, in the process, has benefited from the billions of dollars of investment already made in this area. The technology that is enabling integration, and this new drive toward network flexibility, is already available. The answer lies with the use of silicon as an optical material.

The Successful Platform

Before considering the attributes of silicon as an optical material, it is worth taking some time to define what constitutes a successful and flexible integration platform. Essentially, we are looking for an over-arching technology that will allow us to deliver, through network infrastructure providers, what network operators are demanding.

The most important attribute for a component technology is *functionality*. This is the ability of the chosen platform to perform all (or at least, most) of the tasks (generation, routing, conditioning, and detection of light signals) required by engineers working in optical networking. Such tasks fall broadly into two categories: active functions and passive functions.

These two categories of function can in their turn be realized by two different methods. The first of these is true "intrinsic" integration, which involves building structures within the material of the optical device itself. The chosen material and production system therefore needs to lend itself to this kind of manipulation.

The second method is hybridization—attaching discrete devices directly to the surface of the optical substrate. This technique is valuable because it allows the integration of functions that, for one reason or another, cannot be implemented within the optical material that forms the basis of the device.

Whilst functionality is important in the current context, an integration platform needs several other attributes. Its *performance* will decide whether the resulting components can reach, and if possible exceed, the performance benchmarks required within the industry. Despite economic slowdown, it is important for industry to keep pace with technological development. Indeed, times of economic instability could even encourage the development of faster, stronger, and cheaper components.

Another important feature is *reliability*. This relates both to reliability in-service and to two other key parameters: *manufacturability* and *manufacturing yield.* Components that can be reliably produced by well-understood, repeatable, automated processes will generally produce superior manufacturing yields (that is, less manufacturing defects resulting in rejected parts) and thus be more economically viable. They will also be less likely to fail in the field. Service providers expect to be able to control their networks, and this can boil down to reliability—in theory it should be easier to control and monitor one multifunctional device rather than several linked via fiber.

Service providers also have to deal with the high expectations of their customers, who are unforgiving when faced with a failed system. In the communications industry, reliability is essential to ensure that reputation remains positive.

Physical size and *power consumption* are factors that must also be considered. Smaller components are desirable not just because they take up less space in the final application, but also because with more devices per wafer, production yields can be radically enhanced. Miniaturization also goes hand-in-hand with integration; it is easier to incorporate more functions on each device if the space required for each function is smaller. Additionally, by integrating functions, fewer components are needed, which has a knock-on effect on the interconnections. The fibers connecting various chips take up a huge amount of space, and therefore the number of components used is directly linked to the amount of fiber used. Fewer components lead to a reduction in the amount of fiber, which leads to less bulk and less floor-space cost for the network operator.

Infrastructure vendors are increasingly asking for quick, reliable production of a number of product variants. They want *customized* products that can be sold at different specifications all over the world. In essence, this is another argument for integration.

Bookham Technology's ASOC® technology is a good example of the new breed of component integration technologies which are rising to satisfy these demands.

Based on single-mode rib waveguides (see *Figure 1*) formed on silicon-on-insulator substrates, ASOC uses the fact that silicon is practically "transparent" at telecommunications wavelengths to produce waveguides with low losses (less than 0.2 dB/cm). The refractive index of silicon is high (3.5), allowing compact optical circuits to be made.

The fundamental passive elements of a waveguide-based integration platform include, for example, bends, couplers, and fiber-waveguide interfaces, and these need to be easily combined to produce the desired optical functions, preferably via easily used computer-aided design (CAD) models. A key early stage in the development of ASOC as a successful integration platform was the development of such a set of passive elements that could be assembled into practical integrated optics devices.

The engineering investment in ASOC as a large-scale optical integration platform began to pay off in 1999, as the company began to build a suite of products based on its arrayed waveguide grating (AWG) offering. The AWG is at the core of many dense wavelength division multiplexing (DWDM) functions, performing the key task of combining many wavelengths of light for transmission along a single fiber and splitting those wavelengths at the receiving end.

In volume production, silicon AWGs yield non-adjacent channel crosstalk levels of below –45 dB, enabling worst-case specifications of –23 dB for 40-channel Gaussian devices and eliminating the need for "clean-up" filters. The small-dimension ASOC waveguides were combined with 2D mode expanders to achieve insertion losses of less than 8 dB for flattened devices and less than 6 dB for a Gaussian.

For Bookham Technology, the "building-block" approach has been supplemented by a further step toward higher-level integration within a single device. Doping and injecting charge into silicon waveguides enhances the range of optical functions available through intrinsic integration, rather than by hybridization. Charge injection changes both the refractive index and linear absorption coefficient of the material and allows an electrical signal to be used to attenuate the light traveling through the silicon waveguides. This is a valuable capability since it points directly to other functions commonly used in DWDM networks—signal attenuation and channel balancing—and is a major benefit of using silicon waveguides, as this technique has yet to be proven in other waveguide materials such as silica.

FIGURE 1

The Basic ASOC Waveguide Structure

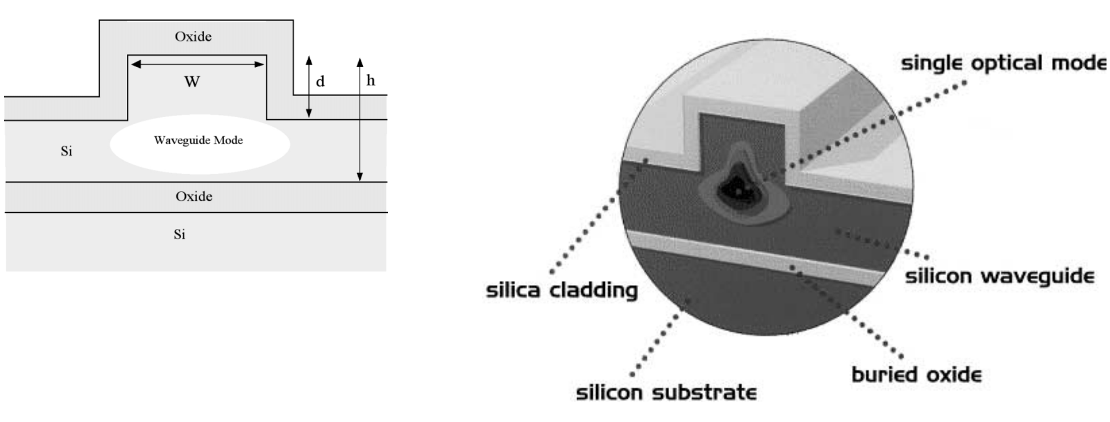

FIGURE 2

Doping and Charge Injection Enables More Complex Functions

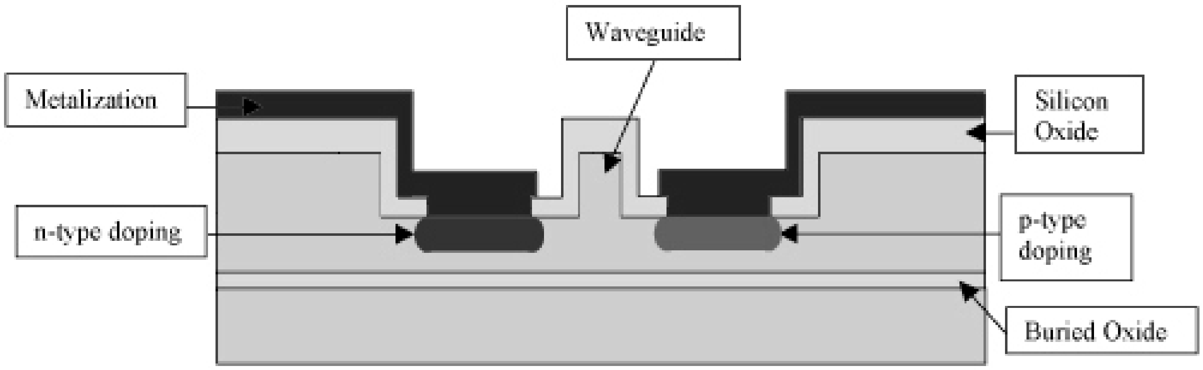

The ASOC absorption attenuator structure (see *Figure 2*) uses a monolithically integrated pin diode to inject both electrons and holes into the waveguide region. The resultant variable optical attenuator devices have a switching speed in the order of megahertz—three orders of magnitude greater than that available in other technologies. By integrating several attenuators on a single chip, Bookham Technology demonstrated an early capability to perform parallel integration (i.e. multiple "copies" of the same function on a single chip) within a silicon platform.

This approach to producing attenuation can also be applied in implementing the switching functions that have always been intrinsic to communications applications. The ASOC attenuator is so fast, and produces of such a high level of attenuation, that it can effectively be used as a 1x1 (on-off) switch—a fundamental building block in communications systems.

We have noted that a wide range of functionality is a key attribute for successful integration and that hybridization is an important capability that radically increases the range of functionality. This was a key initial attraction toward the use of silicon: Mounting other semiconductor elements, such as lasers, on the surface of the silicon device can be done via established processes (see *Figure 3*). Unlike alternatives such as silica, there is little thermal mismatch between the hybridized device and the silicon substrate, and reliability is not compromised. There is also a good match between the refractive indices of the various semiconductor materials, easing the task of coupling light from one material to the other. Lasers, for instance, can be located in etched holes in the silicon, allowing the laser stripe to be positioned against the waveguide facet. Since the waveguide and laser location holes are defined in the same semiconductor etch process, placement accuracy using automated pick and place equipment in high-speed assembly processes is assured.

If we consider a simple end-to-end DWDM system (see *Figure 4*), we can see the importance of this range of functionality. The lasers that are used to generate the light traveling through the system are active devices: The light is multiplexed together using an AWG or similar passive functional block, which is preceded by a set of attenuators that balance the optical power in each wavelength of light. Depending on the span length of the optical link, there may be several optical amplifiers between the sending and receiving end, where a range of demultiplexing, switching, and attenuation functions are required. There is also the need to perform system monitoring, not only to ensure that the system is up and running, but preferably also to provide "trend" information that will allow the network operator to perform predictive maintenance.

Infrastructure makers, under pressure from their customers, want to reduce the complexity of their task. This increases the attractiveness of buying as many functions as possible within an integrated subsystem, in which the various compatibilities and performance trade-offs have already been understood and accounted for.

FIGURE 3

Hybridization Radically Increases the Range of Optical Functions

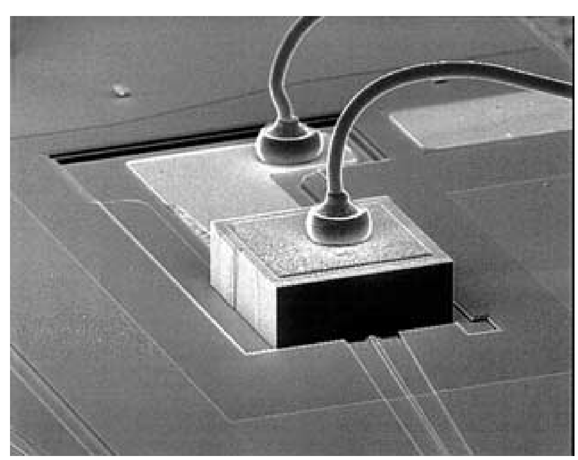

FIGURE 4

Even the Simplest DWDM System Requires a Wide Range of Functionality and Integrated Monitoring

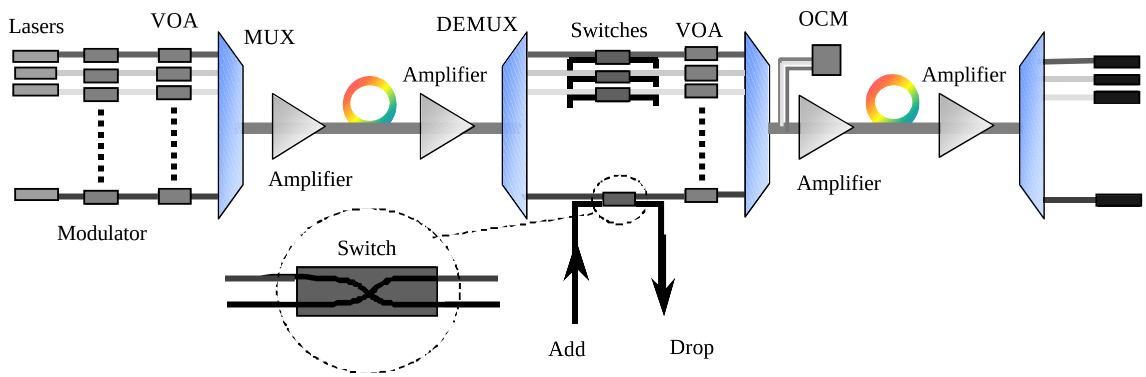

On-chip channel monitoring offers a clear example of how integration can work. Bookham Technology's optical spectrum analysis module (OSAM) with switching input is specifically designed to provide a much more cost-effective way of monitoring power levels in DWDM systems. It reduces the costs of testing from hundreds of dollars per port to tens of dollars. The OSAM checks wavelengths and powers in multiple locations, and this is the best way to cut cost. By putting a switch at the front end of the monitoring system and monitoring different wavelengths sequentially, the need for several traditional channel monitors is integrated into a single chip. The OSAM has successfully demonstrated that increased functionality, switching, demultiplexing and detecting is available in a single package. The market already offers integrated switches, but this is the first time that switching has been fully integrated with other functions on a single chip, without the need to link several tasks via a fiber.

Another example is the integrated multiplexer and variable optical attenuators (Mux–VOAs, see *Figure 5*) and Demux–VOAs, which for the first time are commercially available components that fully integrate the functions of a 40-channel multiplexer/demultiplexer with 40 variable optical attenuators onto a single, compact silicon chip. The list of benefits includes excellent channel uniformity, up to 20 dB attenuation with no degradation of optical performance, very high attenuation accuracy, fast warm-up time, space saving, and low power consumption.

The current environment continues to be difficult, and future visibility remains unclear. However, the challenges posed by the new environment mean that there is light ahead in the optics world, and there will always be a demand for solutions that reduce time and cost and increase network up-time. The opportunities for silicon-on-silicon–based solutions, such as ASOC, will continue. This is a solution with the highest potential to be the industry standard for high-level integration.

FIGURE 5

Mux–VOAs

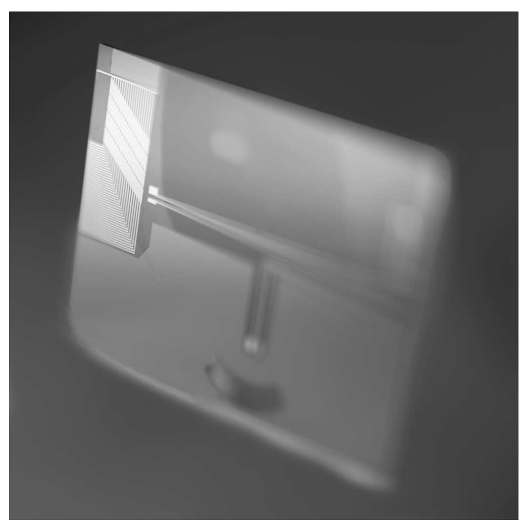

Let's Recount the Votes on All-Optical

Clifford Holliday
President
B & C Consulting

Dense wavelength division multiplexing (DWDM) has brought us the answer to the need for massive amounts of bandwidth across the country, across the world, and now even across the city. The need for this bandwidth has been, and continues to be, driven by the emergence of the Internet (and associated applications) as the single most significant development in the history of communications. This ever-increasing demand has changed the whole approach to facility, and even route, provisioning. Bandwidth increments (e.g., optical carrier [OC]–48) that would have been the entire cross-section on a given route only two or three years ago are now not even provisioning interval increments.

Optical cross-connects (OXCs or optical switches) have been recently recognized as an indispensable component of the evolving optical networks. Some vendors have rushed to claim that the optical version of OXCs (i.e., "transparent" OXCs as opposed to OXCs based on an optical/electrical conversion prior to switching) is the key to the vision of the all-optical networks (AONs). These claims have alleged that only the optical versions of the OXC fit this vision and that the optical/electrical OXCs are somehow inferior.

This article will provide a balanced investigation of OXCs and the advantages and disadvantages of each type. Particularly, we will explore differences as relate to transmission engineering requirements between all-optical OXCs and optical-to-electrical-to-optical (O–E–O) OXCs. As noted, there have been a lot of discussions, claims, and publicity about the "goodness" of an "all-optical" solution. This article is a demand for a recount on some of those claims.

We Had a Bottleneck; Now We Have Gridlock

Some of the carriers (AT&T, for example) are estimating that their backbones are increasing at a rate of 400% a year! While DWDM has given us the tool to deal with the bandwidth explosion, it has not given us a way to be able to control this massive amount of bandwidth. We were looking for a way to keep up with this unending demand for bandwidth, and we found it in DWDM. Trying to keep up with that growth, both the existing growth and the forward looking view of a really scary "hockey stick" growth curve, was the first problem. Now, we can take existing fiber (at least to a great extent—some of the older fiber may not support all of the capabilities of DWDM at all distances), add DWDM, and we have an instant increase in bandwidth of up to 100 times. So, that solves the bottleneck problem, at least for now.

What is left, however, may be an even greater problem. The questions are how to deal with all of this added bandwidth and how to deal with the rapid changes and rearrangements implied by a market growing this fast. In a market growing like this, the end customers are also in turmoil, trying to devise the best way to deal with growth from their viewpoints, thus causing many change orders to bandwidth providers.

In addition to the problems of the raw growth and the high amount of customer churn, the network is largely stuck with an architecture that is singularly unfitted to rapid changes and additions—i.e., synchronous optical network (SONET). To establish a new wavelength on SONET–derived facilities (SONET on DWDM or SONET directly on fiber), it is often necessary to make physical equipment additions and changes at many nodes—particularly where rings intersect. The additions and changes require significant engineering, equipment procuring, site visits, and carefully coordination of testing and turn-up activities.

As a result of all of this work, provisioning intervals stretch into many months and service costs continue to be very high in spite of decreasing costs of underlying infrastructure. These factors have all combined to lead us from a bandwidth bottleneck to a gridlock glut—we have a great deal of potential bandwidth, but it is very hard to effectively use it.

The great national and international carriers have built fiber/DWDM networks that promise and deliver bandwidth that was only dreamed of two or three years ago. Now the issue is learning that quality of service (QoS) involves more than electrical and optical standards. It also involves service flexibility, rapid provisioning and rearrangements, and the ability to support new "on-demand" services. These are the qualities that will support the new strategies of service providers in the era of Internet protocol (IP), the Internet, business-to-business (B2B), etc. They are also the parameters that will turn all of that bandwidth into useful and marketable services.

The answer that is developing to this gridlock glut is the OXC. This new (new in these sizes and with the capabilities now becoming available) device will (1) allow extremely rapid service provisioning and rearrangements; (2) facilitate mesh topologies that more directly support data traffic and that can save on equipment, facilities, and operations costs; (3) provide an extremely fast optical-level circuit (fiber/wavelength/route) restoration capability; and (4) support new services on a wavelength level based on rapid and remotely controlled switching.

What Is an OXC?

An OXC is a device that allows a given number of optical inputs to access a selection of optical outputs. These inputs and outputs may be wavelengths or fibers (with multiple wavelengths). "OXC" is a broad term denoting a class of optical (and opto-electronic, called O–E–O) devices that allow a flexible interconnection of lightwaves. These devices are variously called "optical switches," "optical routers," "optical cross-connects," and a number of combinations of these and proprietary names. In the broadest sense, OXCs replace physical patch panels and provide an optical "point of flexibility."

OXCs are relatively old devices (in terms of Internet time) in that they have been a long-time internal component of most DWDM terminals. In that application, they are very small, simple devices that are connected to detectors (that detect a change in a SONET signal or detect the loss of a light signal) that cause the OXC to switch between fiber paths for restoration or maintenance purposes. The devices and applications now being planned as control centers for our DWDM networks are all, more or less, derivatives of this simple function.

What Is the Difference between Transparent and Opaque OXCs?

Transparent OXC

A transparent OXC (sometimes called an optical or all-optical OXC) has no internal electronics in the lightwave path. It accepts a lightwave as input, reflects the lightwave (using one of a variety of technologies being employed by various vendors), and transmits the lightwave with only the amplitude being changed. Transparent OXCs are controlled by external electronics, but there are no electronic devices in the light path.

Figure 1 shows a particular wavelength (#3) coming in from a DWDM fiber (on the right of the drawing.) That wavelength is converted to a standard 1310 nm wavelength, is switched in the OXC by being reflected, and then is transmitted (on the left) to another DWDM system (channel #5) on another fiber.

Opaque OXC

O–E–O OXCs ("opaque" in the sense that the lightwave does not go through the switch) are the traditional way to perform this function. There are several systems shipping at this point with up to 512 x 512 ports. Many of the systems available are much smaller than this, however, with many limited to 32 x 32.

These switches largely use well-known and widely available application-specific integrated circuit (ASIC) very–large-scale integrated (VLSI) technology. The optical signal is terminated on a diode, where it is converted to an electrical signal. After conversion, the digital information can be monitored, switched, and used as a basis for intelligence in directing the output. It is converted back to an optical signal as the output of a laser.

Figure 2 illustrates an O–E–O OXC in a typical application. It shows a wavelength coming in on a DWDM system (channel #3) operating on a fiber on the right of the drawing. This DWDM wavelength is converted to a standard 1310 nm wavelength, introduced to the OXC, converted in the OXC to an electrical signal, switched in the OXC to an outgoing port, transmitted out as a 1310 nm wavelength, and ultimately coupled to another DWDM system, where it is transmitted out on channel #5 of that system on another fiber.

Comparison of Transparent and Opaque Switches

It should be noted that, although some of the press coverage would seem to suggest otherwise, transparent switches are not inherently better than opaque. Like most other technology choices, they each have their advantages and disadvantages.

FIGURE 1

Example of the Application of Transparent Switch—1310 nm Interfaces

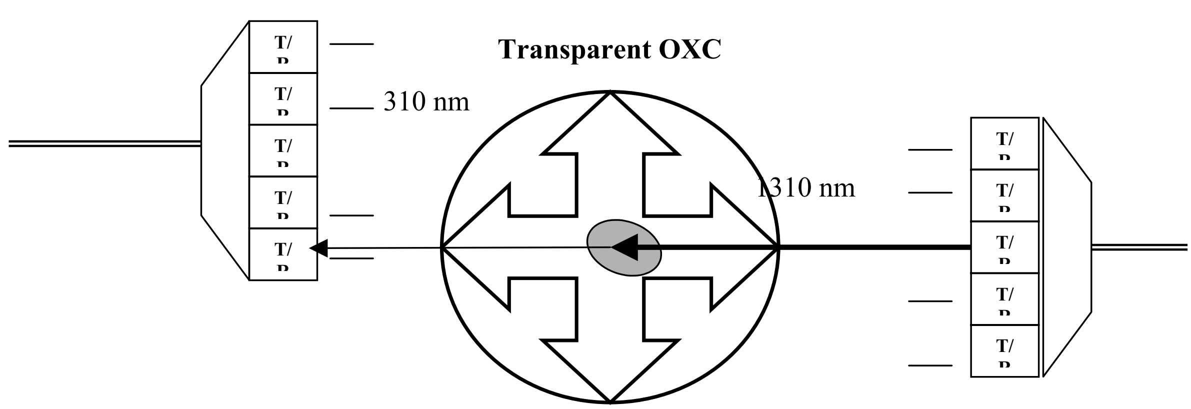

FIGURE 2

Example of an Opaque OXC in a Typical Environment

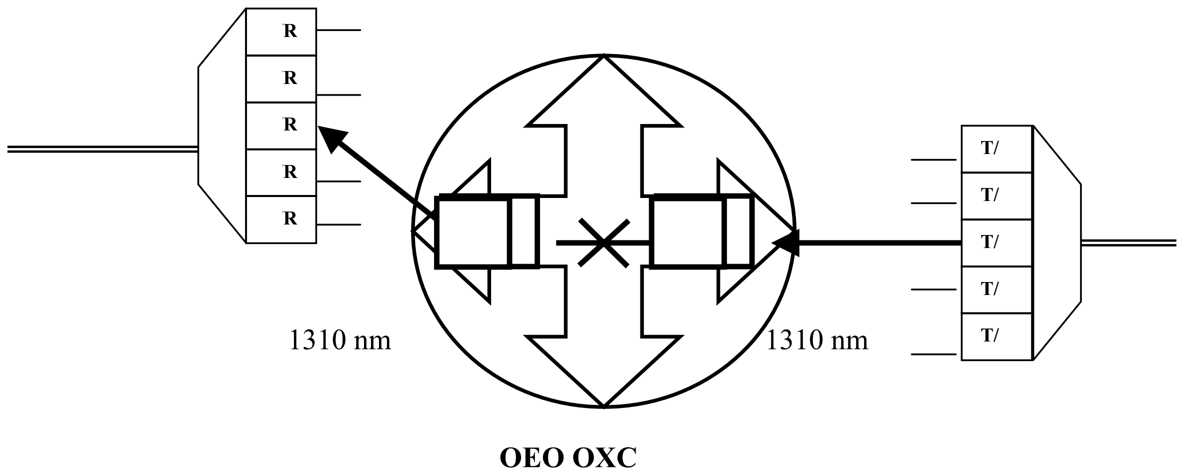

The advantages of the transparent switches are as follows:

- The simple interface and the elimination of the need to buy and deploy the light-to-electrical and electrical-to-light devices.
- The transparent switches are best suited to the vision of the nirvana of the "all-optical network."
- Transparent switches are speed and protocol agnostic, making upgrades easy. Opaque switches are especially bit-rate dependent. This makes upgrades of speed very difficult for opaque switches.
- Transparent switches can probably be built to a larger scale than opaque switches. More than 1,000 ports (in and out each) are clearly going to be shortly available, and many thousands of ports are likely. Opaque switches are going to be limited to 512 ports for at least a while.
- Transparent switches are extremely dumb devices, and thus simple and, at least potentially, cheap.

The disadvantages (at least at this time in the development cycles) are as follows:

- Using transparent devices introduces a loss item into the light-loss budget calculation. The transparent switches are not completely transparent. They reflect light by one means or another, and each reflection point is a loss point. In O–E–O systems, the switches are, in fact, regenerators, and they are improvements in the transmission path rather than detriments.
- There is as yet very little standardization among the makers of DWDM equipment as to the wavelengths that they use. In going through an O–E–O device, this makes little difference, as the right "color" is merely a question of selecting the right laser on the output side. In transparent devices, this is not the case. Light is only reflected through the device, so no change in wavelength can be achieved. The color "in" is the color "out," and that is the only one that can be used in that slot.
- Restoration presents another problem. If a lightwave is split (to establish a standby path), then its power is cut in half. With O–E–O, the split can be made in the electrical domain, and the loss can be easily overcome.
- Performance monitoring is also a problem with transparent systems. There is no real way to make the traditional digital error measurements (or corrections) that are possible in O–E–O systems.
- It is likely going to be 2001 before large-port-size versions of the transparent switches are available in any real commercial quantity. Opaque switches are shipping now.

Example of the Insertion Loss Problem with Transparent Switches

The issue of loss in the use of transparent switches is often misunderstood. The use of transparent switches in "all-optical" networks can easily be (mistakenly) thought of as a transmission advantage, and the press sometimes misleads us to think this is the case. In fact, transparent switches are *the source of optical signal loss* in the calculation of the optical path budget. They all use some kind of reflection to achieve the switching function, and several versions require substantial free space travel of the light beam. Both of these (reflection and free space travel) cause significant loss to the signal. Optical amplifiers can provide the remedy for this problem on a given circuit route. However, when transparent OXCs are introduced in the network, a new problem appears in that network segments may be connected together in a more or less random manner, making the end-to-end transmission design very difficult.

Opaque switches, on the other hand, convert the optical signal to an electrical signal. In the electrical domain, the signal is given the full "3R" (re-amplification, reshaping, and retiming) regeneration treatment, and then it is reproduced as an optical signal at a new, "cleaned-up," high level.

Figure 3 is a sketch illustrating a typical DWDM route with an OXC. This route takes a signal from point A to B on a

FIGURE 3

Normal Route for a DWDM System

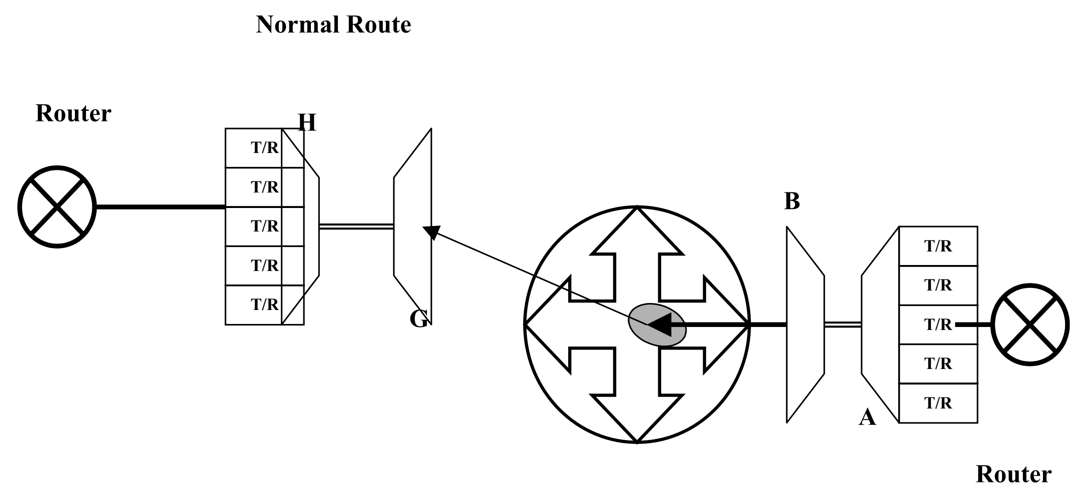

DWDM system. At B, a DWDM terminal interconnects to an OXC and then to another DWDM system (G). This system goes to the DWDM terminal at H and finally to a router located also at H.

Figure 4 shows the situation when a fault has occurred between G and H. The normal path (A–B, G–H) for the light signal through the OXC has been disrupted. The OXC has switched the signal coming from A to a restoration path (C–D, E–F). It also shows this complete hypothesized restoration (A–B, C–D, E–F) path (assuming a fiber fault). The following chart (see *Figure 5*) shows the path-loss budget calculations for the two paths.

The significance of this loss difference depends on the receiver sensitivity and the transmitter output power. If the

FIGURE 4

Example of Power Loss in Transparent Switches

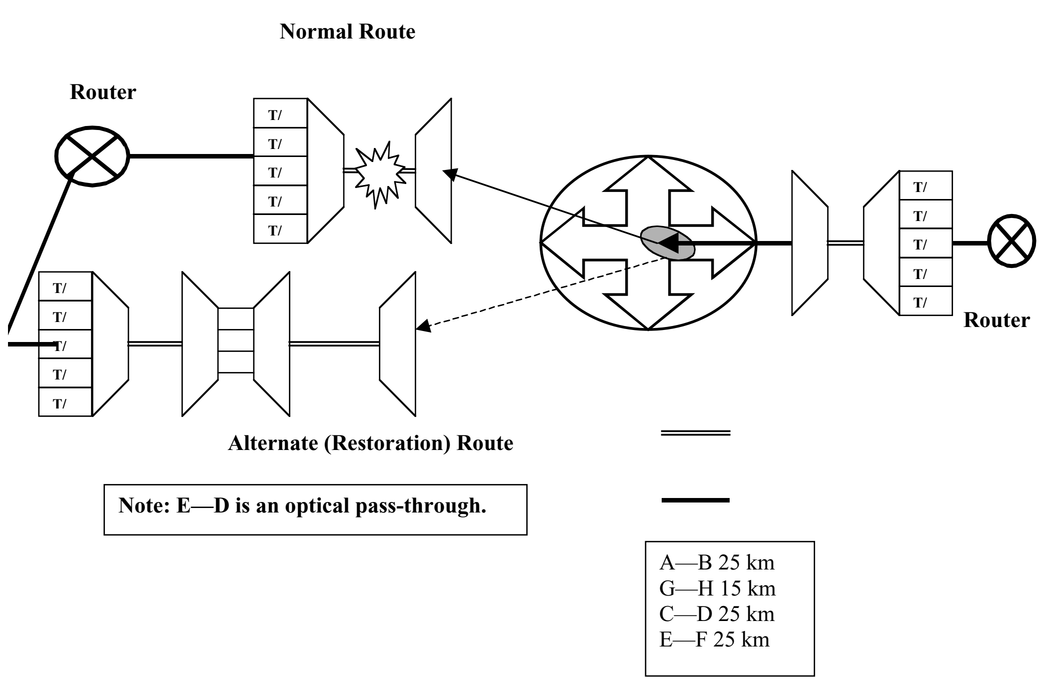

FIGURE 5

Comparison of Loss Budgets—Active and Restoral Paths

	Normal Path (A–B, G–H)	Restoration Path (A–B, C–D, E–F)
Fiber Loss (@0.27/km @ 1550 nm)		
A–B, G–H (40 km)	10.80 db	
A–B, C–D, E–F (75 km)		20.25 db
Connector Loss (@ 1 db)	12 db	14 db
Splice Losses (@ 0.2 db – 1 per 5 km)	1.6 db	2.4 db
Margin for Repair Splices (@ 0.2 db – 2 per section)	0.4 db	0.6 db
Path Totals	**24.8 db**	**37.25 db**
Optical Switch Loss	7 db	7 db
Total Loss	**31.8 db**	**44.25 db**

output power is 0 db, and the sensitivity is –35 db (typical figures, however, some systems would have higher output power), then the normal path would be well within the limits of the receiver (–35 db). However, in the event of a switch for restoration, the alternative path would be out of limits for the receiver, even though the components of that path would individually be in limits. Obviously, the author has concocted this example to illustrate the point. However, the reader should note that all of the assumptions are absolutely real-life situations, and one would expect to encounter this problem in any use of transparent switches.

It should also be noted that if the transparent optical switch were not used (i.e., if an O–E–O switch were used instead), there would be no problem on the alternate path. In that case, the optical-switch loss will no longer apply, and the paths will not be additive for loss calculations. (In other words, the path A–B, C–D, E–F becomes two paths, A–B and C–D, E–F, each of which is within the allowable limits of the receivers. The C–D, E–F loss is 23.1 db, and the A–B loss is much less.)

Thus, there are two insertion loss problems: (1) the combination of routes, which have been individually engineered, and (2) the loss of the transparent optical switch itself. These problems can be overcome by very careful engineering of all conceivable interconnections and by the application of optical amplifiers. Both of these, however, are costly and detract from the main advantages of the all-optical approach—economics and simplicity.

While there are solutions for the transparent OXC problems indicated, the solutions might be worse than the problems. The real answer is that "all-optical" may not necessarily really mean that the OXCs are always fully optical. Both transparent and opaque OXCs have appropriate applications in our new networks. As in any technology choice, one needs to avoid the hype. Rather, potential users should carefully review the available alternatives as they apply their specific situations.

The author's latest "Lightwave" series of reports (www.igi-group.com) is the basis for this article. This article originally appeared in the April 9, 2001, issue of Telephony *as "Solving the Gridlock Glut."*

OSS Challenges for Multiple Broadband Service Integration

Karen Johnson
Chief Information Officer
e-Site, Inc.

Introduction

Technology exists today to support operation support system (OSS) flow-through provisioning. The majority of issues today are not centered around technology constraints, but rather involve business-interaction problems, such as the following:[1]

1. Labor-intensive processes and their associated time and costs
2. Error-prone hand-offs of information between operational entities
3. Delays in provisioning cycles
4. Low levels of service assurance
5. Lack of billing synchronization
6. Multiple proprietary solutions
7. Lack of visibility into partner systems and networks

Standardized flow-through, end-to-end interactions between various trading partners are just beginning to be implemented in broadband integration solutions today.

There are a multitude of broadband forecasts and projections published to date. *Figure 1* offers one such forecast as provided by the International Data Corporation (IDC) in the second quarter of 2000. As individual access provisioning processes among trading partners move to OSS standardization and technology improvements, the question is this: will this affect the rate of consumer adoption of one broadband access technology over another?

The current difficulties with broadband provisioning are highlited by a study conducted by Forrester Research. Participants were asked to identify the worst things about the current telecom purchasing processes. The results are shown in the *Figure 2*. Most, if not all, of these can be solved by the OSS integration of services.

Architecture Design

The architectural design of an OSS solution includes six distinct components, as *Table* 1 outlines.

As companies migrate an OSS solution from one access technology into multiple broadband access technologies, the level of complexity is increased because access technologies remain competitors. OSS standardization efforts are being addressed by each access technology individually.

Multi-Technology OSS Integration

The key to integration is to integrate business logic, not just functional capabilities. This requires consolidating provider-specific workflow and data transformations into a single consolidated software framework.

Legacy systems across technologies differ from, more than they parallel, each other. In addition, partially automated flows within each industry require more complex call-center support. Some industries use a standard version of middleware while others use a proprietary variant. Each provider product, regardless of broadband technology, requires the development of connectors and business logic to drive the connectors.

Application programming interface (API) suites within each industry are underway. However, they do not mimic other broadband technology API suites. The biggest challenge is in the analysis and modeling of data that must support a number of differing technology products used in one integrated solution.

OSS Issues for Multi-Industry Integration

The challenge in creating connectors and business logic is getting cooperation with existing business logic and APIs that vary across industries.

The amount of work required for integration (APIs, data mapping, business-logic definitions, etc.) does not decrease when moving from one industry provider to another industry provider. Fault, configuration, accounting, performance, and security overlap, and the information requirements from one system to another differ. This is especially problematic when the same data is required in two industry-disparate systems.

The biggest issues for multi-technology integration is working through promised API functionality by isolating and

FIGURE 1

IDC's Broadband Deployment Projections

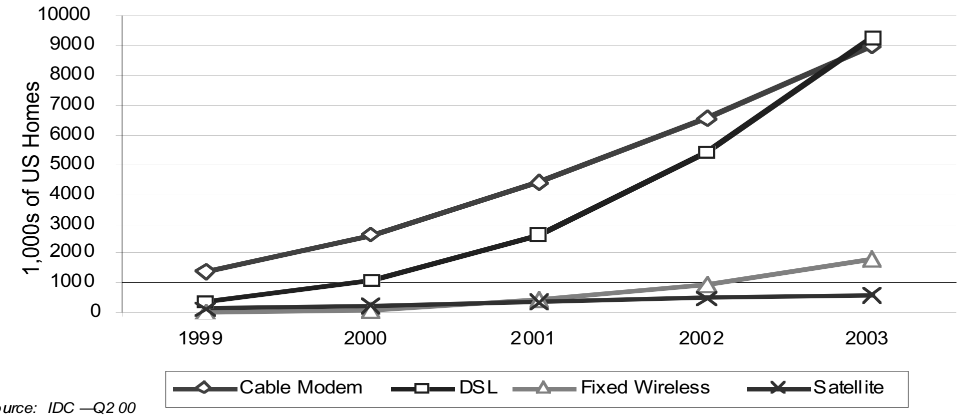

Source: IDC —Q2 00

testing not only the technology-specific data but also the provider-specific data.

The One with the Most Customers Wins

Network and service providers need to work together and across industries to provide customer-driven broadband-based solutions for services, products, and content.

Extending franchises, increasing customer acquisition while up-selling current customers, and creating new revenue streams succeed in any industry.

Next-Generation Broadband Users

Next-generation broadband users are diverse members looking for a broad range of services designed specifically to help them achieve their goals and provide a place where business can be easily transacted. These are the users who want high-speed Internet browsers and are generally quick to add on new services.

Build Community, Not Diversity

Competitors generally do not trust each other. And many broadband companies today are more comfortable going head-to-head rather than working side-by-side. Meanwhile, customers may have to wait several years to have interoperability issues resolved.

However, it is essential that these vendors and providers start working together for long-term viability. Collectively, industries can provide buyers and suppliers with content they cannot do without. Pure economic forces, then, should

FIGURE 2

The Worst Things about Current Telecom Purchasing Processes

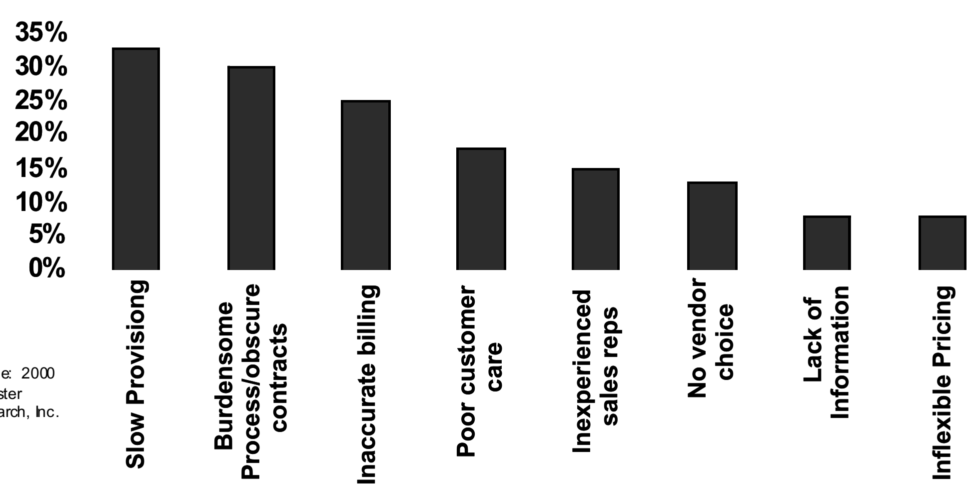

Source: 2000 Forrester Research, Inc.

TABLE 1

Architectural Design Components of an OSS Solution

IT Infrastructure Bandwidth limitations and routing bottlenecks	**Security Policies** Policies for the operating environments	**Data Transport** Data processes and ensured delivery
Data Translations Information for differing OSSs	**Process Mapping** Workflow rules for OSS information flows	**Product Design** Normalized data schema and API conventions

encourage broadband companies to work with, rather than against, each other.

Summary

Successful broadband service providers will be able to do the following:

- Automate business processes
- Integrate back-end systems
- Enable on-line customer acquisition and support
- Extend their platforms into e-business partners and e-marketplaces
- Create more satisfied customers and less churn
- Increase sales, brand trust, and market share
- Adapt quickly to additional revenue sources
- Integrate systems for network and application provisioning
- Adhere to technology standards
- Build community, not diversity
- Build integration architecture that adapts to the evolution of products both within and outside the core technology

Notes

1. *DSL Forum Technical Report TR–038*, March 2000

DWDM Passive Components

Gerd Keiser
President
PhotonicsComm Solutions, Inc.

This paper discusses the passive components that are the enabling technologies for dense wavelength division multiplexing (DWDM). It shows why people are using wavelength multiplexing, and it discusses what international standards groups such as the International Telecommunication Union (ITU) are doing to specify wavelength spacings. Four fundamental DWDM technologies are compared: dielectric thin-film filters, fiber-based Bragg gratings, arrayed waveguide gratings, and bulk diffraction gratings. The discussion covers operating principles, advantages, limitations, and performance characteristics for each of these technologies.

Transmission Windows

The curve in *Figure 1* shows typical attenuation as a function of wavelength for a silica fiber. The curve has a downward trend because of Rayleigh scattering. It also has a large attenuation spike at around 1400 nanometers (nm) because of absorption of the light by water molecules in the glass. Current material processing methods can now eliminate this spike and create a continuous curve, but this paper will assume that the spike is there, as most installed fibers tend to have this characteristic.

The two transmission bands used by telecommunication systems are the so-called second and third windows. Because optical sources have very narrow output spectral widths, many optical sources can fit within these two windows, allowing the simultaneous transmission of many independent information streams at different wavelengths. The spacing between the wavelength channels is specified by Recommendation G.652 from the ITU–Telecommunication Standardization Sector (ITU–T). The standard spacing is 100 gigahertz (GHz), which is equivalent to 0.8 nm at a 1550-nm wavelength. The G.652 document also specifies optional spacings of 25, 50, and 200 GHz.

ITU-T Wavelength Grid for Wavelength Division Multiplexing

Figure 2 shows a typical wavelength grid for 100-GHz channel spacings for 32 channels. This is based on the ITU–T reference frequency of 193.100 terahertz (THz), which is 1552.52 nm in the wavelength domain.

Dielectric Thin-Film Filter

The first DWDM technology discussed here is the use of dielectric thin-film filters (see *Figure 3*). These filters consist of a stack of dielectric thin films separated by cavities. The multiple reflective dielectric thin-film layers form mirrors that surround the cavities, thereby creating a thin-film resonant cavity filter or a Fabry-Perot interferometer. The device is made to pass a particular wavelength straight through and to reflect all others. The length of the cavity determines which wavelength passes through the filter.

As the number of cavities increases, the passband of the filter sharpens up, creating a flat top for the filter, a desirable characteristic. Some advantages of a dielectric thin-film filter are its flat passbands with sharp frequency rolloffs, its low insertion loss (less than three decibels [dB]), and its high reflection loss (more than 50 dB). One of its limitations is that its use for closely spaced channels requires many very uniform thin-film layers for a number of cavities. Also, one filter is needed for each wavelength. Because of these limitations, dielectric thin-film filters are useful for a small number of wavelength channels. Standard commercial devices are available for up to 20 channels with 100-GHz spacings, and some vendors are offering products with smaller channel spacings.

Bragg Grating Principle

A second DWDM technology is based on the Bragg grating principle. These devices are used extensively for functions such as dispersion compensation, stabilizing laser diodes, and add/drop multiplexing in optical systems. One embodiment is to create a Bragg grating in the core of an optical fiber. In such a fiber-based Bragg filter there is a periodic variation in the refractive index along the direction of light propagation. This variation is shown in *Figure 4*, where n_1 is the refractive index of the core of the fiber, n_2 is the index of the cladding, lower-case lambda (l) is the wavelength, and upper-case lambda (L) is the period of the grating. If an incident optical wave at l_0 encounters a periodic variation in refractive index along the direction of propagation, l_0 will be reflected back if the following condition is met: $l_0 = 2n_{eff}L$, where n effective (n_{eff}) is the average weighting of the two indices of refraction, n_1 and n_2. When a specific wavelength l_0 meets this condition, that wavelength will get reflected and all others will pass through.

Fiber-Based Bragg Grating Formation

A fiber-based Bragg grating is typically made from germanium (Ge)-doped silicon fiber, as this material is sensitive to ultraviolet light. One method for forming a Bragg grating in

FIGURE 1

Transmission Windows

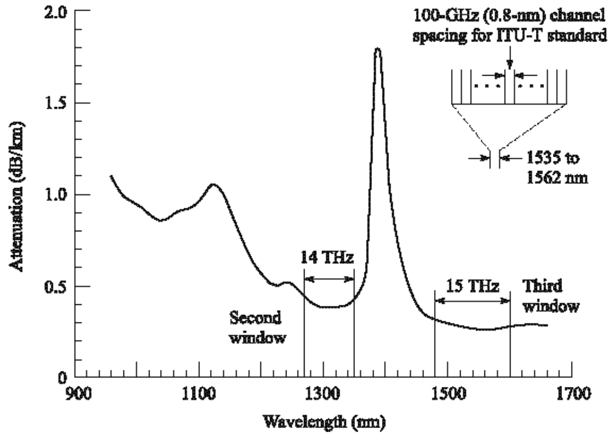

FIGURE 2

32 Channels of the ITU-T Wavelength Grid for WDM Using 100-GHz Channel Spacing

Channel	Wavelength (nm)	Channel	Wavelength (nm)
1	1557.36	17	1544.53
2	1556.55	18	1543.73
3	1555.75	19	1542.94
4	1554.94	20	1542.14
5	1554.13	21	1541.35
6	1553.33	22	1540.56
7	1552.52	23	1539.77
8	1551.72	24	1538.98
9	1550.92	25	1538.19
10	1550.12	26	1537.40
11	1549.31	27	1536.61
12	1548.51	28	1535.82
13	1547.72	29	1535.04
14	1546.92	30	1534.25
15	1546.12	31	1533.47
16	1545.32	32	1532.68

a fiber is shown in *Figure 5*. Two ultraviolet beams at 244 nm are used to establish an interference pattern in the fiber core, which results in a permanent photo imprint of an index variation. The space between these index variations is L (the period of the grating). This fiber by itself is normally temperature sensitive, which will cause the Bragg wavelength to shift with temperature changes. However, special packing techniques can be implemented to compensate for this thermal drift, thus making Bragg gratings athermal devices.

Simple Demux Function

A simple example of how to apply a Bragg grating as a drop element is shown by the demultiplexing (demuxing) function in *Figure 6*. One key element of this process is an optical circulator, which performs the following function: Light that enters port one exits port two, light that enters port two exits port three, and light that enters port three exits from port one. Suppose the Bragg grating is designed to reflect l_2. After the wavelengths l_1 through l_4 enter port one and exit from port two, l_2 gets reflected by the grating, goes back into port two of the circulator, and exits from

FIGURE 3

Dielectric Thin-Film Filter

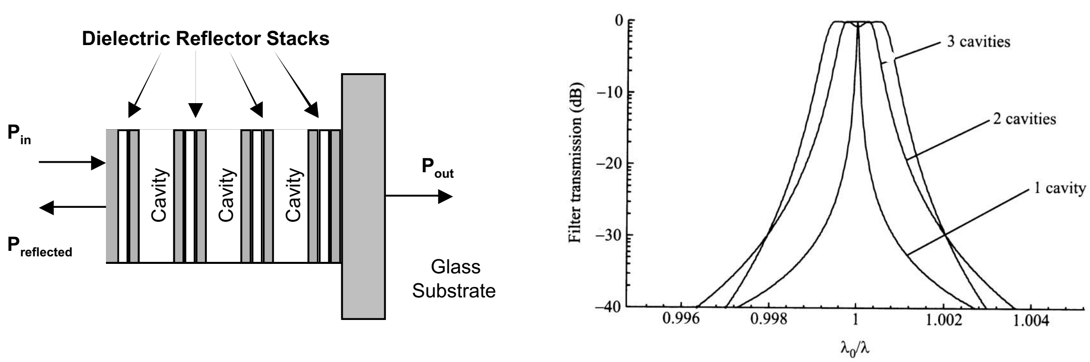

FIGURE 4

Bragg Grating Principle

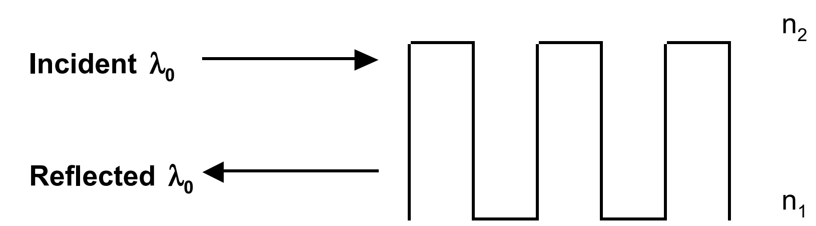

port three. With this procedure that wavelength is dropped and the others pass straight through both the circulator and the grating.

Extended Add/Drop Multiplexing

More complex structures can be made using additional circulators and different gratings in series (see *Figure 7*). These gratings may be made to reflect a fixed wavelength, or they can be tunable to have greater implementation versatility. In this example, there are two three-port circulators and three gratings. The gratings are set to reflect l_1, l_2, and l_3. If a series of wavelengths l_1 through l_N enters port one of the left-hand circulator, l_4 through l_N go straight through both circulators, entering at port one and exiting from port two in each device. The other three wavelengths are reflected by the gratings after the first circulator, re-enter its port two, and exit from port three. These wavelengths then can be demuxed into individual channels and processed. To add these wavelengths back into the fiber-optic stream, they can be sent to a multiplexer (mux). In this case, they would enter port three of the second circulator, exit port one, head toward the first circulator, get reflected by the three gratings, re-enter port one of the second circulator, and exit port two, where they join the other wavelengths.

Fiber-Based Grating Characteristics

Bragg gratings allow wavelength channel spacings as narrow as 25 GHz. By using special packaging techniques, Bragg gratings can be made to have a very low thermal drift, down to less than half a picometer (pm) per degree centigrade, and they exhibit very low interchannel crosstalk. One filter is needed for each wavelength, and normally the operation is sequential, with wavelengths being reflected by one filter after another. Therefore the losses are not uniform from channel to channel, as each wavelength goes through a different number of gratings, each of which adds loss to that channel. This may be acceptable for a small number of channels, but the loss differentials between the first and last reflected wavelengths in a series of Bragg gratings is a limitation for large channel counts.

FIGURE 5

Fiber-Based Bragg Grating Formation

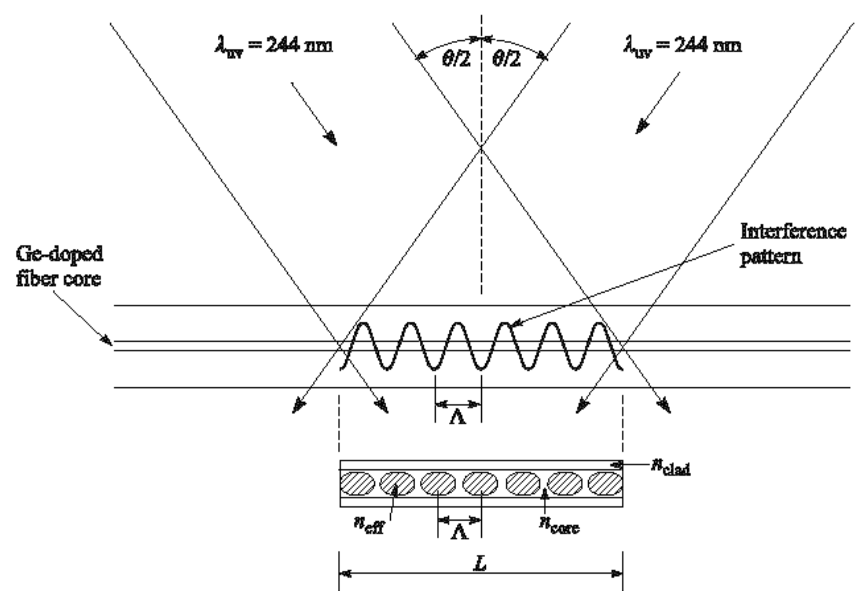

FIGURE 6

Simple Demux Function

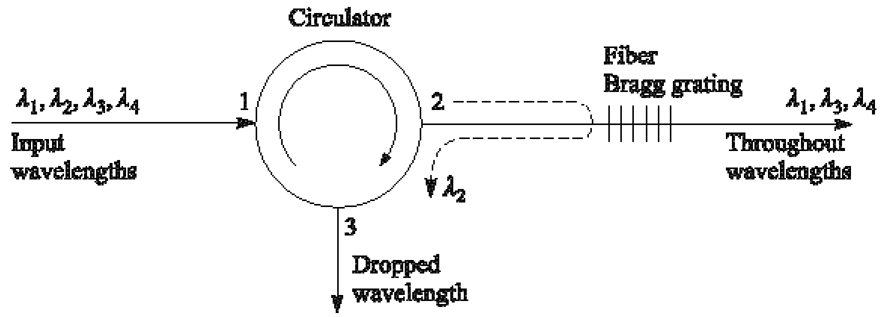

FIGURE 7

Extended Add/Drop Multiplexing

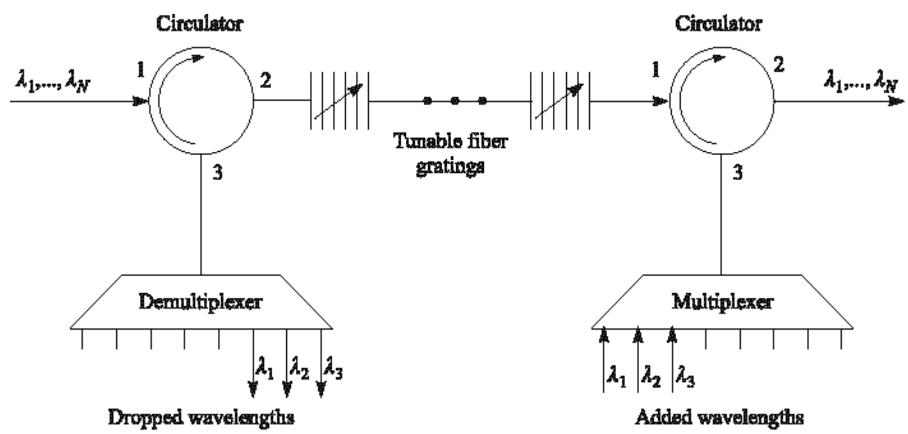

Arrayed Waveguide Grating

Arrayed waveguide gratings (AWG) are a third DWDM device category. These gratings use input and output slab waveguides to focus light in planar star couplers and a series of waveguides with different path lengths called a grating array waveguide. If different wavelengths come in on separate fibers, the path length differences of the grating array waveguides can be chosen such that all input wavelengths exit from one port of the device. The operation is also reciprocal, so that if all the wavelengths come in on one fiber, the device will separate them into individual output channels.

AWG Characteristics

AWG technology makes it easy to achieve very narrow channel spacing of 25 GHz. Devices with 16, 32, and 40 channels are available commercially, and they are easily fabricated on silica wafers. Their compact footprint eases system integration. Originally these devices required a thermoelectric cooler to prevent wavelength drift, but newer packaging designs make the devices athermal so they do not need any special dynamic temperature-stabilizing mechanisms.

Bulk Diffraction Gratings

A fourth DWDM technology is based on bulk diffraction gratings. These gratings are fine-ruled or etched parallel lines on some type of reflective surface. With these gratings, light will bounce off the grating at an angle. The angle at which the light leaves the grating depends on its wavelength, so the reflected light fans out in a spectrum. For DWDM applications, the lines are spaced equally, and each individual wavelength will be reflected at a slightly different angle. There can be a reception fiber at each of the positions where the reflected light gets focused. Individual wavelengths will thus be directed to separate fibers.

Characteristics of Bulk Diffraction Gratings

With bulk diffraction gratings, separating and combining wavelengths is a parallel process, as opposed to the serial process used with the fiber-based Bragg gratings. Adjacent-channel crosstalk is very low, usually less than 30 dB. Insertion loss is also low (less than three dB) and is uniform to within one dB over a large number of channels. *Figure 8* shows ruled-grating insertion loss over 50 channels. The variation between the highest and lowest channels is less than one dB, even for high channel counts, such as 50 channels. A

FIGURE 8

Ruled-Grating Insertion Loss

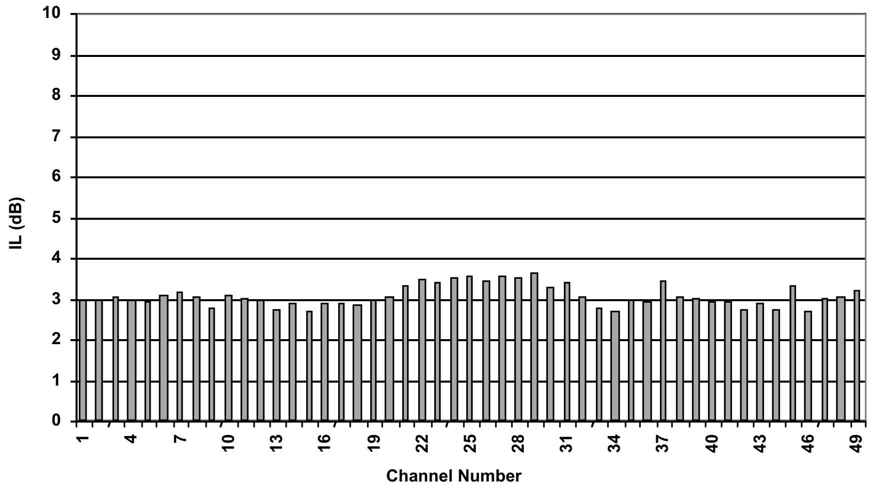

passband of 30 GHz at one-dB ripple is standard. As is the case with other wavelength division multiplexer (WDM) schemes, packaging designs can make the device athermal so that no active temperature control is needed.

Ruled-Grating DWDM Module

Figure 9 shows how this device works. In this schematic, there are several different input fibers. As the light from a fiber spreads out in a cone, there are some simple focusing elements that collimate the light. The lenses are fairly standard components that are readily available. After it is collimated, the light is reflected off the bulk grating. Each wavelength is reflected at a slightly different angle, so that when the light travels back to the array of fibers on the left side of the schematic the optical elements focus each wavelength onto a different output fiber. The device works reciprocally; that is, if different wavelengths come into the device on the individual input fibers, all of the wavelengths will be focused back into one fiber after traveling through the device.

Comparison Summary

Each WDM technology has its own advantages and limitations. *Figure 10* summarizes and compares the four DWDM technologies discussed in this paper. For example, the chart shows that it is difficult for thin-film filters and Bragg grating to handle narrow wavelength spacings. In addi-

FIGURE 9

Ruled-Grating DWDM Module

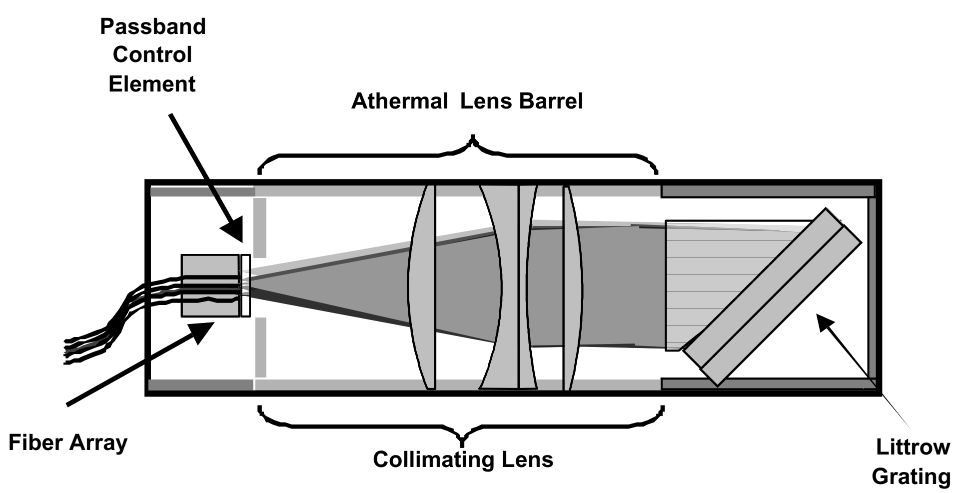

FIGURE 10

Comparison Summary

Parameter	Thin-Film Filter	Fiber Bragg Grating	Arrayed Waveguide Grating (AWG)	Bulk Diffraction Grating
Channel Spacing Accuracy	Good but difficult for < 100 GHz	Good down to at least 50 GHz	No problem down to 25 GHz	No problem down to 25 GHz
Fiber-to-Fiber Loss	Very good; not uniform because of sequential operation	Very good; not uniform because of sequential operation	Uniform; typically < 3 dB overall	Uniform; typically < 3 dB overall
Crosstalk Level	Very good for spacings > 100 GHz	Very good; 30 to 35 dB	Very good; 30 to 35 dB	Very good; 30 to 35 dB
Absolute Wavelength	Each channel is adjusted separately	Each channel is adjusted separately	Good but may need temperature control	Very good; athermal device
Cost per Channel	High	High	Low	Low
Channel Count	Good for < 20 channels at 100 GHz	Good for < 20 channels at 25 GHz	Good for more than 16 channels	Good for more than 16 channels

tion, they work well for a small number of channels, whereas AWG and bulk diffraction gratings are more advantageous for narrow-wavelength spacing and for a large number of channels.

References

1. Boord, W. T., T. L. Vanderwert, and R. DeSalvo. "Bulk Diffraction Gratings Play Increasing Role in Optical Networking," *Lightwave* 18 (March 2001): 172-178.
2. Ehlers, H., M. Biletzke, B. Kuhlow, G. Przyrembel, and U. H. P. Fischer. "Optoelectronic Packaging of AWG Modules and Their Environmental Stability Tests," *Optical Fiber Technology* 6 (Oct. 2000): 344-356.
3. Keiser, G. "A Review of WDM Technology and Applications," *Optical Fiber Technology* 5 (Jan. 1999): 3-39.
4. ———. *Optical Fiber Communications*. New York: McGraw-Hill, 2000.
5. Minasian, R. A. "Photonic Signal Processing of High-Speed Signals Using Fiber Gratings," *Optical Fiber Technology* 6 (April 2000): 91-108.

Comparison of Metro SONET and DWDM

Andrew Knott
Vice President, Marketing
White Rock Networks

SONET and DWDM in the Metro

This paper will compare the metro markets for synchronous optical networks (SONET) and dense wavelength division multiplexing (DWDM). Most vendors seem to take the *Field of Dreams* view of the market: If the vendor builds it, carriers will come. But carriers will come only if it is to their advantage. The success of metro SONET and DWDM will depend on network evolution and carrier spending—not just on the technology.

Background

White Rock Networks, headquartered in Dallas, was founded just over a year ago to provide modular, scalable, optical networking products for the metro market. It is funded by bench capital backed by a start-up, has raised more than $120 million so far, and has more than 200 employees. Its target customers are metro local-exchange carriers (MLEC). Paul Johnson at Robertson Stephens has joked that the *m* in MLEC stands for "monied." The first shipments to these "monied" LECs was to be in August or September of this year, which is important because new vendors in this market are more likely to succeed if they can turn ideas into implementation in a short time to match advances in networks and component technology.

The Carrier Dilemma

In the United States, *a huge growth in bandwidth is projected over the next three to four years* (see *Figure 1*). Similar growth is projected worldwide. There will also be healthy growth in revenue. But the real problem in today's market is that carrier costs are rising faster than revenue. That is one of the reasons why a large number of carriers are now concentrating heavily on providing services where they can actually make money and on getting the best value out of their existing networks. Solutions that provide maximum profits with immediate return on investment (ROI) are crucial.

White Rock's Laws of Carrier Motion

The following laws of carrier motion, in the spirit of Newton's Laws of Motion, condense White Rock's view of the market.

First Law of Carrier Motion

The First Law of Carrier Motion is that every successful vendor-carrier relationship tends to remain successful unless an external force is applied. That external force could be a competitor with a product that does much the same thing for less money, or it could be the vendor trying to do something that the customer does not want.

Second Law of Carrier Motion

The Second Law of Carrier Motion is that force is equal to money times acceleration (F = m*a), where force is the amount of force that a vendor uses to convince carriers to adopt a new method of operation; money is, in most cases, revenue; and acceleration is the rate at which that money will move from the vendor to his competitors. When a vendor shows a new technology to a large carrier, the response the vendor dreads most is that the technology is interesting and the carrier is studying it. "Studying" in carrier-speak means "do not expect to see revenue from this product in your lifetime."

Third Law of Carrier Motion

The Third Law of Carrier Motion is that for every disruptive vendor push, there is an equal and opposite carrier push. If a vendor attempts to do something that a carrier is not ready for, the carrier will tend to push back—and push to other vendors.

The Laws of Carrier Motion Applied

The laws of carrier motion can be applied much the way a high school physics teacher might give students a set of rules and then see how they apply to the real world.

When Vendors Push, Carriers Push Back

A good example of the Third Law—when vendors push, carriers push back—is the way radical technology changes tend to force disruptive operational cost changes that can outweigh the short-term cost benefits of the new technology. This is one reason why a lot of new technology does not get absorbed into large carriers, or gets absorbed only through fairly slow adoption cycles. For example, two to three years ago, every industry report predicted that asynchronous transfer mode (ATM) over SONET would be the technology that was used in the

FIGURE 1

The Carrier Dilemma

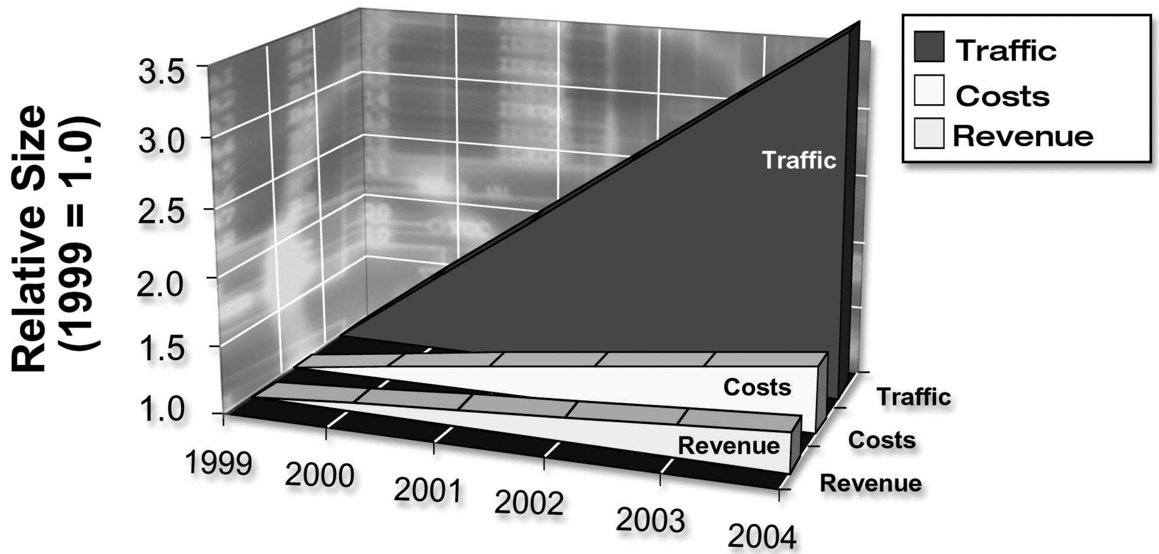

network, and that Internet protocol (IP) and other technologies would run over ATM over SONET. But currently, there appear to be no successful ATM over SONET products.

When Carriers Pull, Responsive Vendors Stay Synchronized

Things really happen in the network when carriers pull and responsive vendors stay in synchronized motion with their customers, following the first law. Staying synchronized is a challenge for the vendor, however, because carriers can normally predict what their networks will look like in the next year, but not in three or four years. That can be a real problem if the vendor's product development cycle is three to four years.

Force Equals Money Times Acceleration

New products that allow immediate cost reductions without forcing radical changes in existing network architecture tend to break in quickly. In other words, most products that involve network evolution rather than revolution succeed extremely well (the Second Law). Cerent's 454 product, otherwise known as the Cisco 15454, is a good example of this in the transmission world. Prior to Cerent's introduction, no new SONET vendor, other than the original five, had produced revenues much over $30 million a year.

Metro Market Size and GrowthæLaws of Motion in Action

The metro market reflects the deployment of SONET products that evolve networks, rather than disrupt networks (see *Figure* 2). A chart from Ryan Hankin Kent such as that in *Figure 2*, showed the last stitch of SONET being sold in North America next year. However, the numbers in *Figure 2* show that SONET remains the largest market in North America and has the highest projected dollar growth. Also, the DWDM market is growing and will represent about 30 percent of the overall metro market by 2004, but it will not completely overwhelm the network. SONET and DWDM will coexist, because carrier networks will force them to.

Metro Market Environment

Diverse Application Space

The metro market environment is very different from what most vendors would like it to be. It is a diverse application space that requires multiservice capabilities, while most vendors would find it easier to design products for customers that look the same. Site-by-site selectability is also needed, because a small competitive local-exchange carrier (CLEC) in a single city is very different from WorldCom or Southwestern Bell. And the metro market requirements in Louisville, Kentucky, are very different from those in San Francisco. While there is a requirement for multiservice delivery, there is a bigger requirement to be able to customize that delivery cost-effectively. Carriers love cheap products. It is also necessary to be able to fit the application to the kind of building that is being entered.

Technology Changes Quickly, but Carriers Do Not

Metro technology changes extremely rapidly. Test component vendors, replace their component technologies every nine to 12 months. So one way for a vendor to stay ahead is to continually refresh and build next-generation technology into its products.

But carriers want the latest technologies only to the extent that they can reduce prices, allowing them to generate fast profits while using their existing networks. They want the benefits of the new technology, but they cannot change architectures in their network every nine to twelve months. Carriers adopt new architectures slowly, but they want immediate cost benefits over present mode of operation (PMO).

FIGURE 2

Metro Market Size and Growth—Laws of Motion in Action

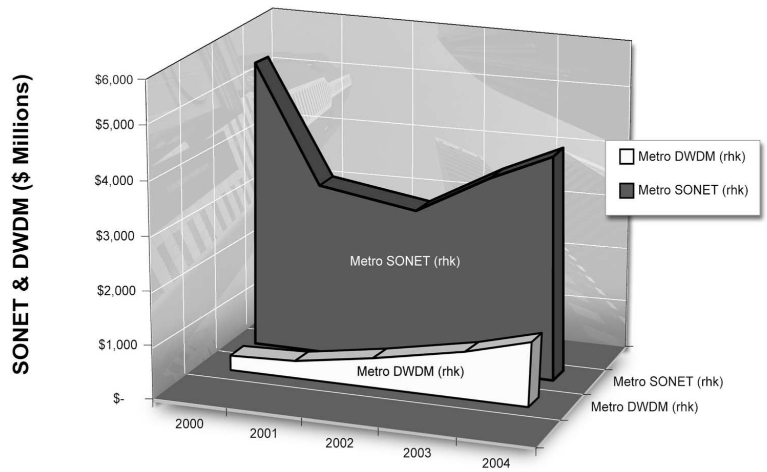

	2000	2001	2002	2003	2004
□ Metro DWDM (rhk)	$401	$549	$717	$973	$1289
■ Metro SONET (rhk)	$5074	$2882	$2441	$3195	$3768

Fiber Deployment Is NOT Homogeneous

Fiber deployment is anything but homogeneous in the United States. About 10 percent of the buildings in the United States are served directly by fiber. The other 90 percent are served by copper or wireless. A typical large CLEC, interexchange carrier (IXC), or incumbent local-exchange carrier (ILEC) has either *no* fiber-to-the-building (FTTB) or an *excess* of FTTB, with multiple strands not being used.

Metro SONET and DWDM Compared

SONET

SONET has the great advantage of having been around for 10 years or so. It is a highly reliable transport for high-bandwidth services. And SONET has been able to grow with the growth of bandwidth in the United States. As recently as three years ago, the default SONET speed used in a metro environment was optical carrier (OC)–3. This year it is OC–48, and next year 30 to 40 percent of deployments will use an OC–192 line rate.

SONET is also ubiquitous. It has well-defined operations that are understood by the vast majority of carriers, and they have operational practices to control it. It is very good at grooming and performance monitoring (PM) at both the transport level and the end-to-end service delivery level. And it enables efficient and extremely non-complex transport of digital signal (DS)–3, OC-n, and data services.

The introduction of low-cost OC–48 add/drop multiplexers (ADM) like the Cerent 454, which came with local area network (LAN) service delivery capabilities, vastly increased the viability of SONET–based service delivery. Next year, there will be volume deployment of low-cost OC–192 ADMs that are metro-capable. There may be a 40 to 50 percent fall in the price of short-haul OC–192 SONET products over the course of the next year. That would be the same factor that pushed OC–48 into the metro market. And the introduction of Gigabit Ethernet (GbE) interfaces onto those products will stretch the life of SONET in the metro.

So SONET addresses the bulk of today's metro market needs and the applications that the majority of "monied" LECs deploy today. The most notable exceptions are sporadic fiber relief and storage area networks (SAN).

DWDM

DWDM can be broken down into two categories: metro fiber relief and wavelength services.

Metro Fiber Relief

Three aspects of DWDM provide metro fiber relief. First, there are DWDM solutions with very low cost, allowing for immediate return on investment (ROI). Second, these solutions are simple and easily understood. And third, it is relatively straightforward to design simple passive DWDM networks, eliminating power and management complexities. A proposed standard has been filed by Lucent with the International Telecommunications Union (ITU) for coarse wavelength division multiplexing (CWDM). With existing DWDM technologies, it is possible to multiply vastly the amount of service that can be deployed in a network without really having to learn too much.

Wavelength Services

Wavelength services are still awaiting prime-time feasibility outside of greenfield carriers. There is a fundamental need for lower component costs and SONET–like management processes in this arena. Also, these products typically provide no direct support for DS–3, which is still viewed as the bread-and-butter service solution by most large carriers. For example, the latest reports show that DS–3 service at BellSouth grew by 60 percent year-over-year.

Also, wavelength services waste fiber potential. The cost of wavelength services simply cannot be justified today for individual service delivery below OC–48, and few customers want individual services at OC–48 and above. A better option, therefore, is a low-cost time division multiplex (TDM) aggregator in front of the DWDM terminal, but only where necessary. That allows multilevel grooming capabilities and PM visibility, and it allows maximum use of fiber potential for the wavelengths that have been deployed.

SONET and DWDM Provide Complementary Benefits
SONET and DWDM provide complementary benefits. SONET provides grooming capabilities and project management at the transport and end-to-end services levels, and its management capabilities are well defined and understood. It also offers low-cost, direct support of DS–3, OC–n, and GbE.

DWDM can be used in a spot fashion for fiber relief, or for direct support of services that are not SONET–compatible, such as fiber channel or services that operate above OC–48 line rates. These capabilities are part of existing carrier networks today, and any technology that is introduced in the future will have to continue to provide these services.

Why Not Just Integrate SONET into Next-Generation DWDM Platforms?

Cost
Integrating SONET into next-generation DWDM platforms is an appealing concept when trying to set up a company, but the biggest reason not to do it is cost. Only a small proportion of locations in an access network require both SONET and DWDM capabilities. And if SONET and DWDM are integrated, carriers will pay a tax for embedded features whether they are using them or not. Also, interoperation will be required between devices that do both DWDM and SONET and devices that do only DWDM or only SONET. Few carriers will voluntarily support two different types of TDM (SONET) multiplexing (integrated and standalone) in a single customer access ring in a single city. "Over-integration" results in a much slower time to market and a sub-optimal cost performance. It is a vendor "push" against carriers' love for low cost.

Complexity
The other reason not to integrate SONET and DWDM is complexity. Carriers have tamed the issues raised by deploying DWDM or SONET alone, but they have not yet tamed the issues raised by deploying both of them on the same product at the same time, including the impacts on management, employees, training, and operations.

Metro SONET and DWDM—A "Dis-Integrated" Solution

A cost-feasible solution that addresses the bulk of today's metro market needs is to deploy SONET and DWDM together in a dis-integrated fashion (see *Figure 3*). In this solution, DWDM or CWDM provides additional wavelengths on a spot basis where carriers start to run out of fiber. SONET is used for the majority of the multiplexing techniques and supports the majority of the customer interfaces. SONET plus CWDM provides very low equipment cost and up to eight lambdas (l) per fiber. SONET plus DWDM provides low equipment cost and better fiber usage.

White Rock Networks offers a building-block family of metro products. Every product is pure play, so no single product has DWDM and SONET in the same platform. Customers can clip products together like Lego blocks to build up exactly what they want to use in a metro environment. The architecture allows products to be built with a 12- to 15-month development cycle, so they can be synchronized with the planning cycles of customers and with the development cycles of suppliers (component manufacturers). The products can thus piggyback on rapid enhancements in optics and optic electronics technology.

White Rock Networks applies the Dell model to the transmission space: technology is continually refreshed to

FIGURE 3
Metro SONET and DWDM—A "Dis-Integrated" Solution

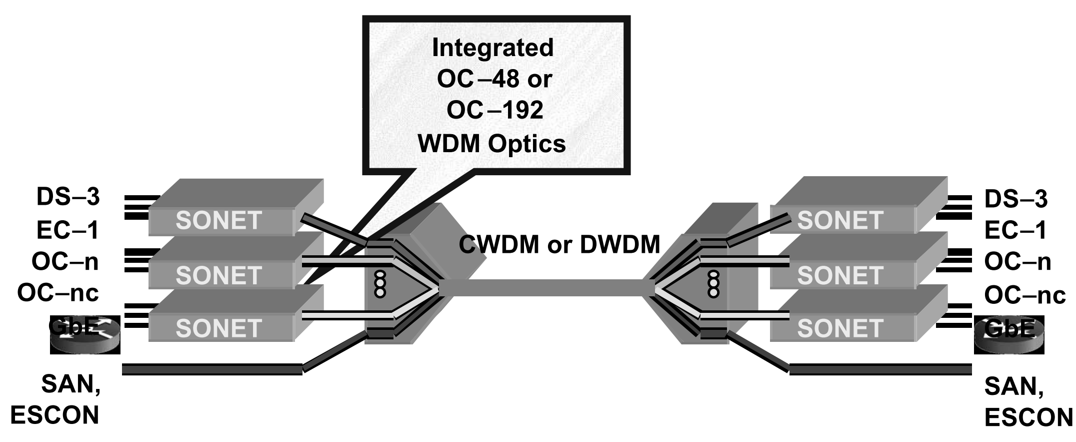

bring best price points and functionality to customers. However, selling carriers a whole load of discrete boxes, even if they are cheap and small, could make carriers' management problems worse. So all the products come with control-playing technology, so as carriers clip them together, they form a single network element. White Rock Networks creates evolutionary optical networks with revolutionary technologies. The first product, announced less than three months ago, is a SONET ADM that is one-twentieth the size of the most common SONET ADM shipped today.

Summary

DWDM can offer significant benefits in fiber consolidation and support for services with very high bit rates, greater than OC–48. SONET, however, is ideally suited for integrated grooming, delivery of individual services up to the OC–48 per customer line rate, and deterministic service protection. So far, large carriers have been slow to standardize "all-optical" solutions. It is the combination of SONET and DWDM in a discrete fashion that will remain the large-carrier access model for the foreseeable future.

Generalized Objected-Oriented Simulation Frame Work for Multimedia-on-Demand Services

Anup Kumar
Professor and Director, Mobile Computing Lab (MINDS)
University of Louisville

Alok Srivatava
Director, Web Services Development
Oracle Corporation

Adel S. Elmaghraby
Acting Chair and Professor, CECS Department
University of Louisville

K. A. Kamel
Dean, College of Information Technology
United Arab Emirates University

Abstract

In today's information age, a growing need is felt for multimedia-on-demand services to provide a variety of facilities such as movie-on-demand, interactive network games, remote education, catalogue browsing, multimedia libraries, home shopping, etc. Currently these facilities are not available widely on a large scale. In this paper, we present a simulation of a generic system to support multimedia-on-demand services over an asynchronous transfer mode (ATM) network. The simulation has been performed using an object-oriented simulation (OOS) approach. Discrete event system specification (DEVS) methodology has been used in the design and implementation of the simulation. This simulation approach also helped in the development of a modular, reusable, and easily extendable model of a multimedia-on-demand system. Also, this approach is shown to be very useful in adapting the basic model to simulate any of the aforementioned multimedia-on-demand services. To illustrate this, we have simulated a movie-on-demand service using our generic multimedia-on-demand simulation model. We have demonstrated the real-time nature, scalability, and cost-effectiveness of a movie-on-demand system based upon our generic multimedia-on-demand system. This model will be useful for studying the performance of services such as movie-on-demand, interactive network game, and multimedia libraries.

1. Introduction

The advent of fast networks, large storage capacity, and high computing power has not only served to make multimedia-on-demand services feasible but has also triggered a strong interest in the development of applications in this area. Basically, multimedia-on-demand services refer to the services that provide user-controlled viewing of multimedia (audio, video, text, graphics, animation, etc.) information. These services include facilities such as movie-on-demand, interactive network games, remote education, catalogue browsing, multimedia libraries, home shopping, etc. Several systems have been proposed that provide a subset of the aforementioned facilities [6,7]. But most of these systems are not very widely available on a commercial basis.

A multimedia-on-demand system has to be designed keeping in mind its feasibility (i.e., whether the current state of technology can support the designed system) and economic viability and performance (i.e., whether the system can meet the required constraints). The existence of high-speed networks such as broadband integrated services digital networks (B–ISDNs), which use ATM, and rapid advances in storage and very–large-scale integrated (VLSI) technology have made multimedia-on-demand systems feasible. The cost and the performance of such a system are heavily influenced by the storage scheme used to store the multimedia data [7,5] type of network and the quality of display. Therefore, it is important that these factors be carefully analyzed before building the system. Three widely accepted techniques that are used in design evaluation to ensure economic viability and performance are prototyping, analytical studies, and simulation.

Building a prototype of a multimedia-on-demand system is expensive because it requires specialized software and hardware. Moreover, complete evaluation of all of the services will not be possible with a small prototype. In general, analytical studies are best for evaluating the performance of a system, but they provide mostly the bounds (upper or lower) on the performance. The main complexity in analytical study comes from 1) the average system performance analysis, 2) sensitivity analysis of various inputs on the system performance, and 3) the interactive nature of an application. Therefore, simulation is the best method to analyze the average system performance and study the sensitivity of various inputs on the performance of the multimedia-on-demand system.

In this paper, we have presented an object-oriented model [1] for a generic multimedia-on-demand system using the C++ programming language. Design details of this system can be found in [5]. Since we are concerned with time-related behavior of the system, a discrete event-simulation approach has been used. To formally present the model and the simulation technique of our system, Zeigler's DEVS methodology [2,3,4] has been adopted. This methodology helps in the hierarchical and modular development of a simulation model of complex systems. The modular nature of our model helps in simulating different storage methods, network parameters, and variable quality displays with minimal changes. To illustrate the generality of our multimedia-on-demand system model, we have simulated a movie-on-demand system and also have compared the performance of the movie-on-demand system for different storage methods and network parameters.

This paper is organized as follows. Section 2 explains the DEVS formalism. Section 3 explains the system architecture in detail. The simulation model is discussed in Section 4. Simulation details are given in Section 5. The simulation studies performed and the results obtained are discussed in Section 6. Section 7 concludes the paper.

2. DEVS Formalism

DEVS is a set-theoretic formalism [2,3] that supports hierarchical and modular specifications of discrete event models. Semantics of this formalism is highly compatible with object-oriented specification of a simulation model. Having set up a simulation model in DEVS, its implementation using an object-oriented programming language (such as C++, MODSIM, etc.) is straightforward. The object-oriented design also supports the distributed implementation of the simulation models. DEVS provides a method to represent a basic model as an atomic unit. Using these atomic units, larger models can be built. A method to lump and/or prune the models has also been proposed by Ziegler [4], which helps in the simplification of simulation models.

2.1 *Atomic Models in DEVS Formalism*

The atomic model, is described as follows:

$M = < X, S, Y, \delta_{int}, \delta_{ext'}\ \lambda, ta >$ where:
X: Input event set
Y: Output event set
S: Sequential states set
δ_{int}: $S \rightarrow S$, Internal transition function
δ_{ext}: $QxS \rightarrow S$, External transition function
λ: Output function
ta: $S \rightarrow R$, Time advance function
$Q = \{(s,e) \mid s \in S, 0 \leq e \leq, ta(s)\}$: Total state of M

In this model, input event set (X) is a collection of all the inputs that influence this atomic model. Information about the input event set is important in order to decide the internal and external transitions δ_{int} and δ_{ext}. The output event set (Y) decides the influence of this model on other models (atomic or coupled). The sequential state set (S) describes the internal states of the atomic model that can be reached by a specified input event from a specified state. Depending upon the current state of the atomic model and the transition rules δ_{int} and δ_{ext}, an output is generated. On an internal event, the output is produced by the output function (k) depending upon the current state of the atomic model. The time advance function (ta) describes the time in which an external event is scheduled to occur. If an external event arrives within ta(s), where s is the current state, then δ_{ext} decides the transition, otherwise δ_{int} decides the transition. The total state (Q) of an atomic model M represents the current state (s) of that component and the time elapsed (e) in that state. A detailed description of the specification of atomic model could be found in [4].

An atomic model is used to define every indivisible component of a simulation model. Using atomic models, other models called molecular models are developed, which have a specified behavior after their composition is defined. Molecular models can also be classified as coupled models. A formal method to represent coupled models in DEVS is given next.

2.2 *Coupled Models in DEVS Formalism*

Several atomic models can be coupled to form a coupled model. This model can be used as a component in a larger coupled model. Thus, complex models can be constructed in a hierarchical fashion. A coupled model so generated has the advantage of information hiding (known as encapsulation in object-oriented methodology). This feature of coupled models lends itself to be considered as an object. A coupled model (CM) is defined as the following:

$CM = < D, (M_I), (I_{ij}), (Z_{ij}), SELECT >$, where:
D: Component names set for each component i in D
M_i: DEVS model for component i in D
I_i: Set of influences of i for each element in I_i
Z_{ij}: Output transition function for transition from from component i to j
SELECT: Subsets of D $\rightarrow$ D, tie-breaking selector

In the previous definition of a coupled model, D is a set of the components upon which this model is constructed. These components could be an atomic model or another coupled model. M_i is a DEVS model for each component that is a part of the coupled model. This information is necessary to allow the internal or external events to be processed by an appropriate component.

3. Architecture of Multimedia-on-Demand Systems

Our multimedia-on-demand system is comprised of three major components (see *Figure 1*), namely the storage server, display equipment, and a fast interconnection network. *Figure 1* explains the organization of the complete system. A brief description of the major components of the system and their responsibilities is given next.

FIGURE 1

Components of a Multimedia-on-Demand System

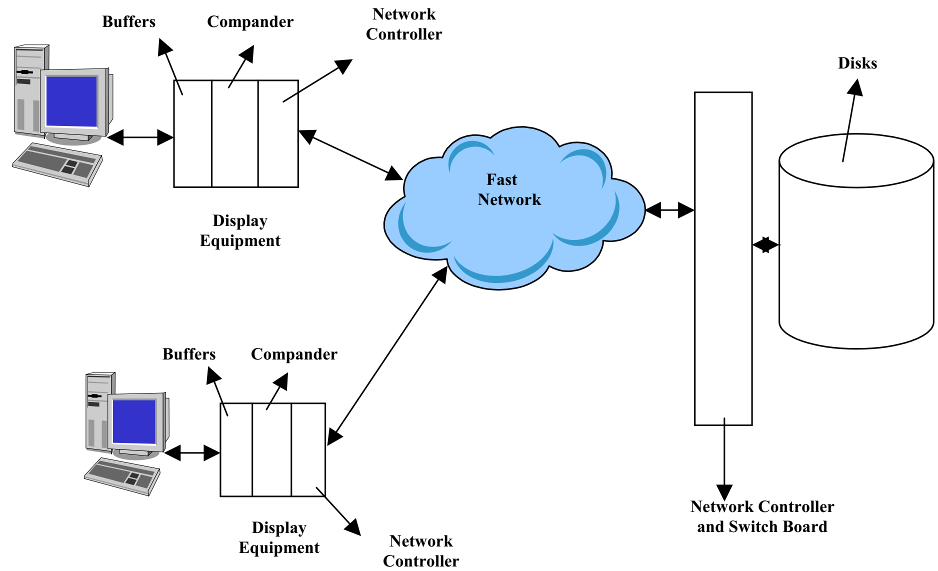

3.1 Storage Server

The storage server is responsible for controlling access to the multimedia data stored in its disks and for transmitting the requested information to the users. The storage server comprises a fast computer, a large number of disks, a broadband switch (e.g., ATM switch), and an efficient network controller. The storage server accepts requests from the users and finds out the appropriate switch (which gets data from one of the disks in the disk array) that can satisfy the request. Then it transmits the incoming information from that switch to the user. The storage server can multicast the same data to a large number of users from the same switch if they happen to have requested the same information at the same time.

3.2 Display Equipment

Display equipment, in our system, is an intelligent display unit capable of sending/receiving signals and multimedia data to/from the server. The display unit can perform real-time uncompression of media information. Thus, a compander, RAM, a network interface, and a high-resolution display constitute display equipment. Each display-equipment unit also has a remote control and/or a joystick for interactive viewing.

3.3 Fast Network

The interconnection network is assumed to have a high bandwidth, guaranteed sequential delivery of packets, and bounded delay in packet-delivery B–ISDN networks using ATM are best suited for providing multimedia-on-demand services. Detailed information on B–ISDN and ATM can be found in [8,9].

4. Simulation Model

We have developed a simulation model of our multimedia-on-demand system using DEVS methodology. The simulation model, shown in *Figure 2*, is a hierarchical coupled model and is object-oriented in nature. The major components of our models are a storage server, display equipment, a global clock, a coordinator, a simulation engine, and a statistics collector. Among these components, the storage server and the display equipment are coupled models, and the rest of them are atomic models. Coupling of these components has been performed according to the discussion in Section 2.2. Due to space limitation, only the details of coupled models have been given. The atomic models can be constructed easily by the description of the component and the methodology given in Section 3.1 [5].

The model of the overall system can be represented using DEVS methodology as follows:

Multimedia-on-Demand System = <

(storage server, display equipment, global clock, coordinator, network, simulation engine, statistics collector),

{{M_{network}}, {CM_{server}} ($CM_{\text{display equipment}}$), ($M_{\text{clock}}$}, {$M_{\text{coordinator}}$} ($M_{\text{simulation engine}}$), ($M_{\text{statistics}}$)),

$Z_{\text{network to server}}$ = Send user request (e.g. some remote control button pressed or joystick moved etc.),
$Z_{\text{network to display equipment}}$ = Send multimedia unit or control signal,
$Z_{\text{server to statistics}}$ = Update statistics,
$Z_{\text{server to network}}$ = Send user request,
$Z_{\text{display equipment to network}}$ = Send user request,
$Z_{\text{display equipment to statistics}}$ = Update statistics,
$Z_{\text{clock to server}}$ = Provide current time,
$Z_{\text{clock to display equipment}}$ = Provide current time,
$Z_{\text{clock to coordinator}}$ = Provide current time,
$Z_{\text{coordinator to server}}$ = Request to begin processing,
$Z_{\text{coordinator to display equipment}}$ = Request to begin processing,
$Z_{\text{coordinator to simulation engine}}$ = Inform when simulation ends,
$Z_{\text{simulation engine to coordinator}}$ = Request start of simulation,
(The rest of the interactions are null, hence they are not given.)}.

Select >

Where Select resolves a tie by priority.

As can be seen from the formalism of the model, the emphasis is only on the interaction of the various components with each other. The internal details of the components are hidden,

FIGURE 2

Model of the Multimedia-on-Demand System

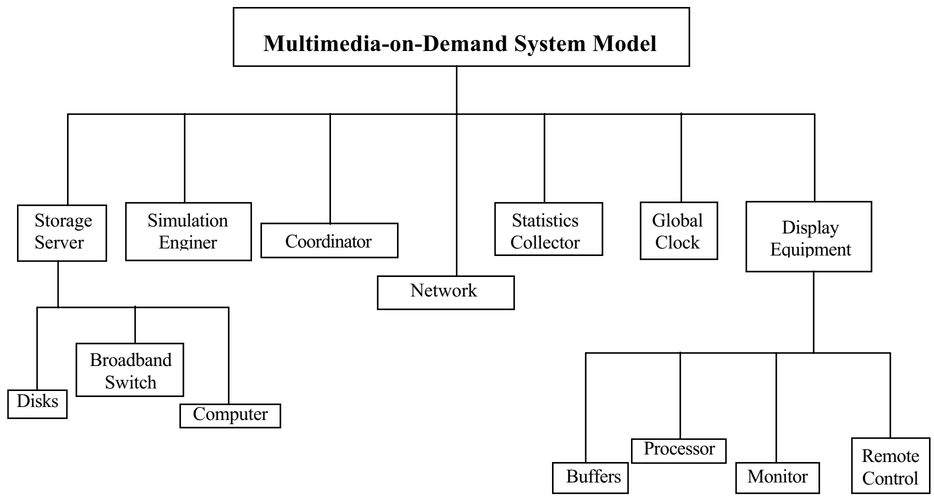

which helps in making the system general (i.e., even if the internals of any component are changed, the model will be unaffected as long as the interactions among the various components are not changed). This feature of the model made it possible to simulate and study the performance of a multimedia-on-demand system for different storage methods and network parameters. A brief description of the components of our simulation model is given in the following sections.

4.1 Model of Storage Server

The storage server comprises an array of disks, a broadband switch, a network controller, and a computer. The server receives inputs from the simulation coordinator and the display equipment and sends the requested data to the display equipment. The server's interaction with the display equipment takes places via the network controller. Since a multimedia-on-demand system is highly interactive, a computer is required to process all the user requests. The server generates statistics about the search and access time for the desired data.

Several constrained allocation schemes have been proposed [5,7] for the storage of multimedia data. Depending on the type of on-demand application, an appropriate layout scheme has to be chosen. In our model for disk (discussed in Section 4.1.1), the model of our storage server can be represented using DEVS methodology as follows:

Storage server = <{array of disks, network controller, broadband switches, computer},

{{M_{disks}}, {$M_{network\ controller}$}, {$M_{broadbases\_switch}$}, {$M_{computer}$ },

$Z_{disk\ to\ broadband\ switches}$ = In every update cycle the data is made available at the corresponding broadband switch buffer,
$Z_{broadband\ switches\ to\ network\ controller}$ = Transmit data to the respective display equipment,
$Z_{network\ controller\ to\ computer}$ = Send user requests to computer for processing,
(The rest of the interactions are null, hence they are not given.)},

Select >

…where Select first accepts the requests coming from the display equipment before the server performs any other operation and $M_{<unit\ name>}$ represents the atomic model for that unit. A brief discussion on the aforementioned atomic units is given next.

4.1.1 Array of Disks

An array of disks is used to store the multimedia data. The layout of this data is discussed in [5]. The disks are responsible for updating their current head position at appropriate intervals of time (decided by the disk retrieval rate) and sending the currently accessed multimedia unit to the server on request. The parameters guiding disk operations in our simulation are disk retrieval rate and disk seek/latency times.

4.1.2 Broadband Switches

In our model, the broadband switches are assumed to be high-speed switches with negligible switching delays. A small number of buffers are attached to each switch to take care of processing delays. The retrieved multimedia units are placed in the buffers by the disks. The switches are responsible for transmitting the multimedia units from their buffers to the appropriate display equipments.

4.1.3 Computer

An intelligent processor is required to maintain the status of each display equipment unit, to handle the interaction between the server and the display equipment, and to locate an appropriate disk for satisfying a user request. For our simulation model, we have assumed that processing time is negligible. This does not cause any inaccuracy in our simulation model because the processing time for any of the aforementioned tasks is at least an order of magnitude smaller than the disk retrieval time, communication delays, etc.

4.1.4 Network Controller

The network controller provides an interface between the storage server and the interconnection network. The computer transmits multimedia data/control signals to and receives control signals from the display equipment via the network controller over the interconnection network.

4.2 Fast Network

The interconnection network provides a communication path between the server and the display equipment. The interconnection model has been designed such that it can incorporate different values for network parameters. Using our model, we can simulate several networks, such as transmission control protocol (TCP)/IP, ISDN, and B–ISDN. In our simulation model, the interconnection network causes a bounded delay in each transmission. The network delay has a significant influence on the buffer requirement of the display equipment. As the network delay increases, the network can also lose packets with a certain probability.

4.3 Model of Display Equipment

The display equipment of a network controller includes a set of buffers, a compander, a monitor, a processor, and a remote control. The display equipment receives input from the coordinator, the server, and the remote control. A processor processes the input coming from the remote control. The input from the server is received via the network controller. The display equipment is responsible for transmitting all user interaction requests to the server. It also collects and transmits statistics about display starvation, start-up delays, buffer overflows, etc. to the statistics collector. To simplify the model, the compander has been lumped with the monitor. Using DEVS formalism, the display equipment model can be expressed as the following:

Display equipment = <
{network controller, buffers, monitor, remote control, processor},
{{$M_{\text{network controller}}$}, {$M_{\text{buffers}}$}, {$M_{\text{monitor}}$}, {$M_{\text{remote control}}$}, {$M_{\text{processor}}$}},

{$Z_{\text{network controller to buffers}}$ = Put an available multimedia unit in the buffer,
$Z_{\text{buffer to monitor}}$ = Transmit a multimedia unit whenever monitor requests,
$Z_{\text{remote control to processor}}$ = Send the user request to the processor.
$Z_{\text{processor to network controller}}$ = Send the user request to the network to be transmitted to the server,

(The rest of the interaction are null, hence they are not given.)},

Select >

…where Select resolves a conflict (if there is one) between an outgoing remote control signal and an incoming server data/signal. A conflict may arise if the two signals arrive at the same time. $M_{\text{<unit name>}}$ represents the atomic model for that unit. The atomic units, which have not been discussed in the context of the server, are described next.

4.3.1 Buffers

Every display-equipment unit has a number of buffers that store the data to be displayed. This is needed to maintain a jitter-free display. Since the number of buffers is limited, it is possible that an overflow may take place and that some data may be lost. The optimal number of buffers required depends on the type of multimedia-on-demand service. One of the studies performed in our simulation model is to determine the optimal number of buffers for a jitter-free display.

4.3.2 Monitor

The monitor performs the functions of decoding multimedia units and displaying them. The time taken for decoding depends on the coding technique used. As soon as the monitor finishes decoding a multimedia unit, it examines the buffers to get the next multimedia unit. If the buffers are empty, then the monitor records it as a jitter.

4.3.3 Remote Control

The remote control provides the interactive nature to the multimedia-on-demand system. The actual, physical form of a remote control may vary depending upon the kind of on-demand service (e.g., it may be a conventional remote control in movie-on-demand systems, a joystick in an interactive network game, etc.), but conceptually it has the same meaning in all services. The remote control transmits user signals to the display equipment's processor, which in turn transmits it to the server. In our simulation model, the remote control generates signals randomly to emulate user actions.

4.3.4 Processor

The processor provides an interface between the remote control and the server. It coordinates the display at the monitor with the remote-control actions (e.g., when a "Pause" button is pressed in a movie-on-demand system, the processor freezes the display on the monitor and at the same time informs the server of its action.) In our model, the processing delays are assumed to be negligible.

4.3.5 Network Controller

The network controller provides an interface between the display equipment and the interconnection network. The processor receives multimedia data/control signals and transmits control signals via the network controller over the interconnection network.

4.4 Global Clock

We have used a global clock in our simulation, which controls the local time management of all of the elements of our simulation model. The clock tick for this clock is set to the highest common factor of the time units used in the components of our model. This clock tick ensures that any event in our system will occur only at multiples of the global clock tick. The actual value of this clock tick depends upon the type of the on-demand application as well as on the parameters of the atomic units. For example, in *Figure 3* the relationship among the time unit associated with disks, monitor, and the global clock has been shown. It can be seen that although the time units for disks, monitor, and the global clock are all different, they are still synchronized.

4.5 Coordinator

In our simulation model, the coordinator is an atomic model that is responsible for coordinating the message exchange between the server and the display equipment. The coordinator initiates the global clock tick and updates it during the simulation. But it does not update any of the local clocks (the local clocks are updated by their "owner" units). The main function of the coordinator is to initiate the system operation and manage the flow of control between the server and the display equipment. To simulate concurrent events on a sequential machine, the coordinator allows all of the components to generate an event at every global clock tick. If a conflict occurs because of multiple events being generated at the same time, the appropriate tie-breaker (referred to as Select in DEVS terminology) is used to resolve the conflict. Coordinator also instructs all of the components of our simulation model to collect and transmit statistics to the statistics collection module. When the simulation time is over, the coordinator dies after returning control to the simulation engine

FIGURE 3

An Example of Clock Synchronization

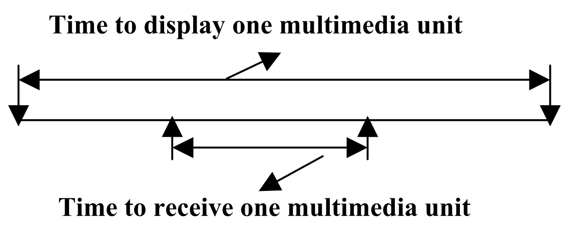

4.6 Simulation Engine

The simulation engine has the responsibility of interpreting the simulation data and of presenting it to the user in accordance with the goals of the simulation exercise. This is a very critical module, and any error in interpreting the data can lead to unrealistic results about the system. If the error is not very obvious, the performance of the simulated system may get erroneously evaluated. In our simulation model, the simulation engine accepts all of the user-defined parameters and requests the coordinator to begin the operation of the system. At the end of the simulation, when the control is returned to the simulation engine, it queries the statistics collection module for the simulation statistics and performs any necessary translations before presenting it to the user.

4.7 Statistics Collector

Statistics collection is vital to any simulation, as the primary purpose of performing a simulation is to collect meaningful statistics about the system's behavior. Thus, the statistics collector module is very important and needs to be designed carefully so that the desired information about the system under study is collected. In our simulation, the statistics collector is basically a data container for statistics pertaining to each module. It is passive in the sense that the responsibility of updating the statistics lays with the respective modules and not with the statistics collector. Also, it does not generate any events of its own and does not have a local clock. The statistics collector module is initialized with the proper values before the simulation is actually started by the simulation engine.

The complete organization of our simulation model is given in *Figure 4*.

5. Simulation Detail

Using our generic multimedia-on-demand simulation model, we have simulated a movie-on-demand (MOD) system. A brief discussion on a MOD system is given next.

5.1 MOD System

The aim of a MOD system is to provide facilities similar to a video cassette player (VCP) while serving as a video rental store. Essentially, MOD can be thought of as an electronic video rental system. The primary advantage of MOD is that it delivers customer-selected movies at a customer-specified schedule. Thus, unlike a conventional VCP, the customer does not have to go to a video rental store. Also, a movie can be selected at the customer's will (unlike in pay-per-view [PPV] services). The MOD service is targeted toward residential customers seeking video information for educational or entertainment purposes.

Market surveys indicate that there is a great demand for MOD services. Although several systems have been proposed in the recent past [6,7], none of the systems provide all of the facilities of a VCP and, moreover, cannot support a large number of users at a reasonable cost. We have proposed

FIGURE 4

Interaction of All the Components in the Simulation Model

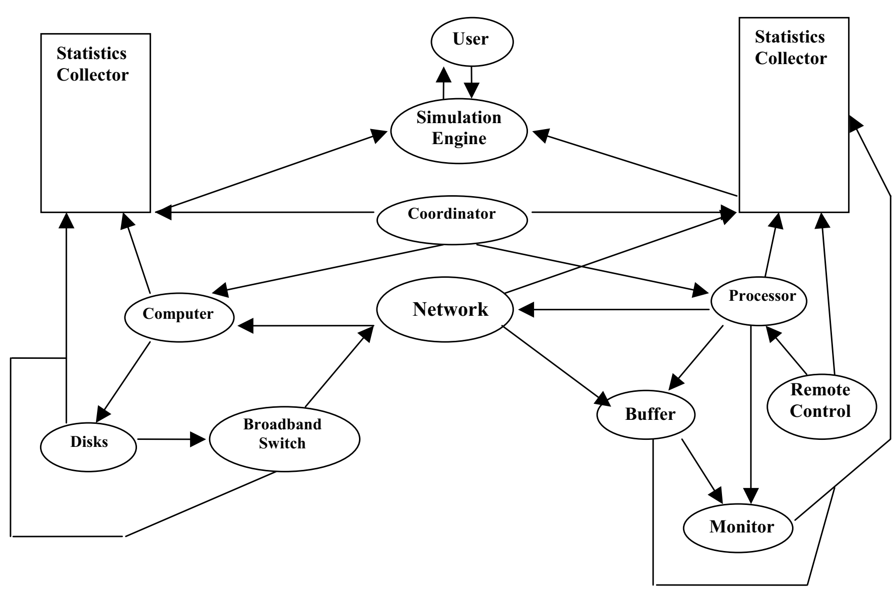

an MOD system that is cost-effective and can support a large number of users while providing all the features of a VCP in [8]. In our MOD system, the server stores a number of movies in an array of disks. Further, multiple disks are used for each movie. Since relatively inexpensive disks are used to store a movie, the use of multiple disks does not increase the overall cost of the system appreciably. The disks are divided into normal and special disks depending on their usage. When display equipment is in rewind or fast-forward mode, it is transmitting data from one of the special disks. When it is in normal mode, it is transmitting data from the normal disks. In our design of the MOD system, the use of multiple disks serves to enhance performance in terms of reducing certain delays (e.g., the delay after which a user starts viewing the requested movie, delay in activating the remote-control features, etc.) in the system. In our scheme, increasing the number of disks can reduce these delays. This is supported by the simulation studies presented in this paper. Also, our scheme is scalable with respect to the number of users in the system, i.e., the performance of the system is not affected as the number of users in the system is increased. The details of our scheme can be found in [5].

In this paper, we present simulation studies of our MOD system. Two different data layout schemes are examined for our MOD system, and their effect on various system performance parameters is compared.

5.2 Parameters Used in the Simulation of a MOD System

In our simulation, we have considered National Television Standards Committee (NTSC) quality video for all movies. The display rate is assumed to be the same for all of the display equipments. The display rate used in our simulation is 1.5 Mbps. All of the movies served by the server are 100 minutes long. Multiple disks with a data retrieval rate of 100 Mbps are used to store the movies. The data unit size of a movie has been kept at 0.05 Mbits (with a 30-frames-per-second display rate, this is the size of a frame of a movie). In our simulation, the display equipment can generate random requests for the *fast forwarding or rewinding with display, fast forwarding or rewinding without display, Pause*, and *Start (Play)*.

Since network delays are assumed to be bounded, their effect on system performance is studied by varying the lower and upper bounds on the network delay. The number of customers in the system is also kept variable.

5.3 Simulation Strategy

In this section, the simulation dynamics are discussed in detail. *Figure 5* gives the complete flow of control of simulation of MOD. Since our simulation is time driven, it is essential to relate the events precisely with the time. In our simulation, all events are time stamped and the delays experienced by an event are also kept with the event. This method ensures that an event is never processed before it reaches the destination. To allow concurrent events, each event generated is stamped with its birth time (decided by the statistical method chosen to generate events; in our simulation, we have generated all events randomly), and its arrival time at every module is also recorded. This allows more than one event to have the same time stamp. In our simulation, at any tick of global clock, all modules are visited in hierarchical order. If the global clock tick matches with the clock tick of the module being visited, an event can be generated. Otherwise, the module does not do anything and passes on the control. As shown in *Figure 5*, the MOD system starts by updating the clock, and then it updates the disks, i.e., accesses a multimedia unit (if available) and feeds it to the broadband switch for transmission. Then the server is asked to check if there are any unattended requests pending to be considered at the current clock tick. If there are some pending requests, the server services them first. Then it checks whether some of the display equipments need a multimedia unit. If so, the respective broadband switch is instructed to send the multimedia unit to that display device. The multimedia unit transmitted carries the time when it was transmitted and the delays experienced, if any. This ensures that, when control is passed to the display equipment, it does not process a media unit, which is physically there but logically should not have reached. Now the display equipment is allowed to process the pending server replies for its requests generated in past. The display equipments also display a multimedia unit at this time (provided there is one available and the display is not stopped because of some remote-control action). Next, the display equipments are requested to generate any requests for the server. If a request is generated, it is transmitted to the server for its consideration at a later time. This completes one simulation cycle. If the simulation time is not yet over, the clock is updated and the next simulation cycle begins. During a simulation cycle, the modules also update the statistics collector. At the end of the simulation, the maximum, minimum, and average values of the collected statistical data are displayed to the user. The analysis of this data is left to the user.

6. Simulation Results

As explained in Section 5.2, the simulation studies are performed for MOD services over an ATM network. The simulation model can be used to compare different layout schemes and network parameters with respect to the following performance measures: 1) average delay in starting a requested movie (referred to as the startup delay), 2) average delay in a remote-control feature (Pause/FF/REV) to become active, 3) average delay in resuming normal display after the remote-control feature has been inactivated, 4) buffers required at the display equipments, 5) average number of jitters experienced in the display, 6) maximum number of display equipments that can be supported by the MOD system, and 7) optimal number of normal and special disks required per movie.

In the simulation studies presented in this paper, we have compared contiguous and constrained (developed in [8]) allocation schemes for our MOD system with respect to the buffers required at the display equipments. The start-up delay and the average delay in the activation of a remote-control feature for the constrained allocation scheme are also discussed.

6.1 Buffers Required at the Display Equipments

The buffer requirements at the display equipments depend heavily on the data layout scheme employed and the network delay. A comparison of the constrained and contiguous data layout schemes with respect to buffer requirement is presented in *Table 1*. Each cell of the table shows two values for the buffer requirement for a particular network delay and the amount of movie passed. In

FIGURE 5

Flow Control in the Simulation Model of the MOD System

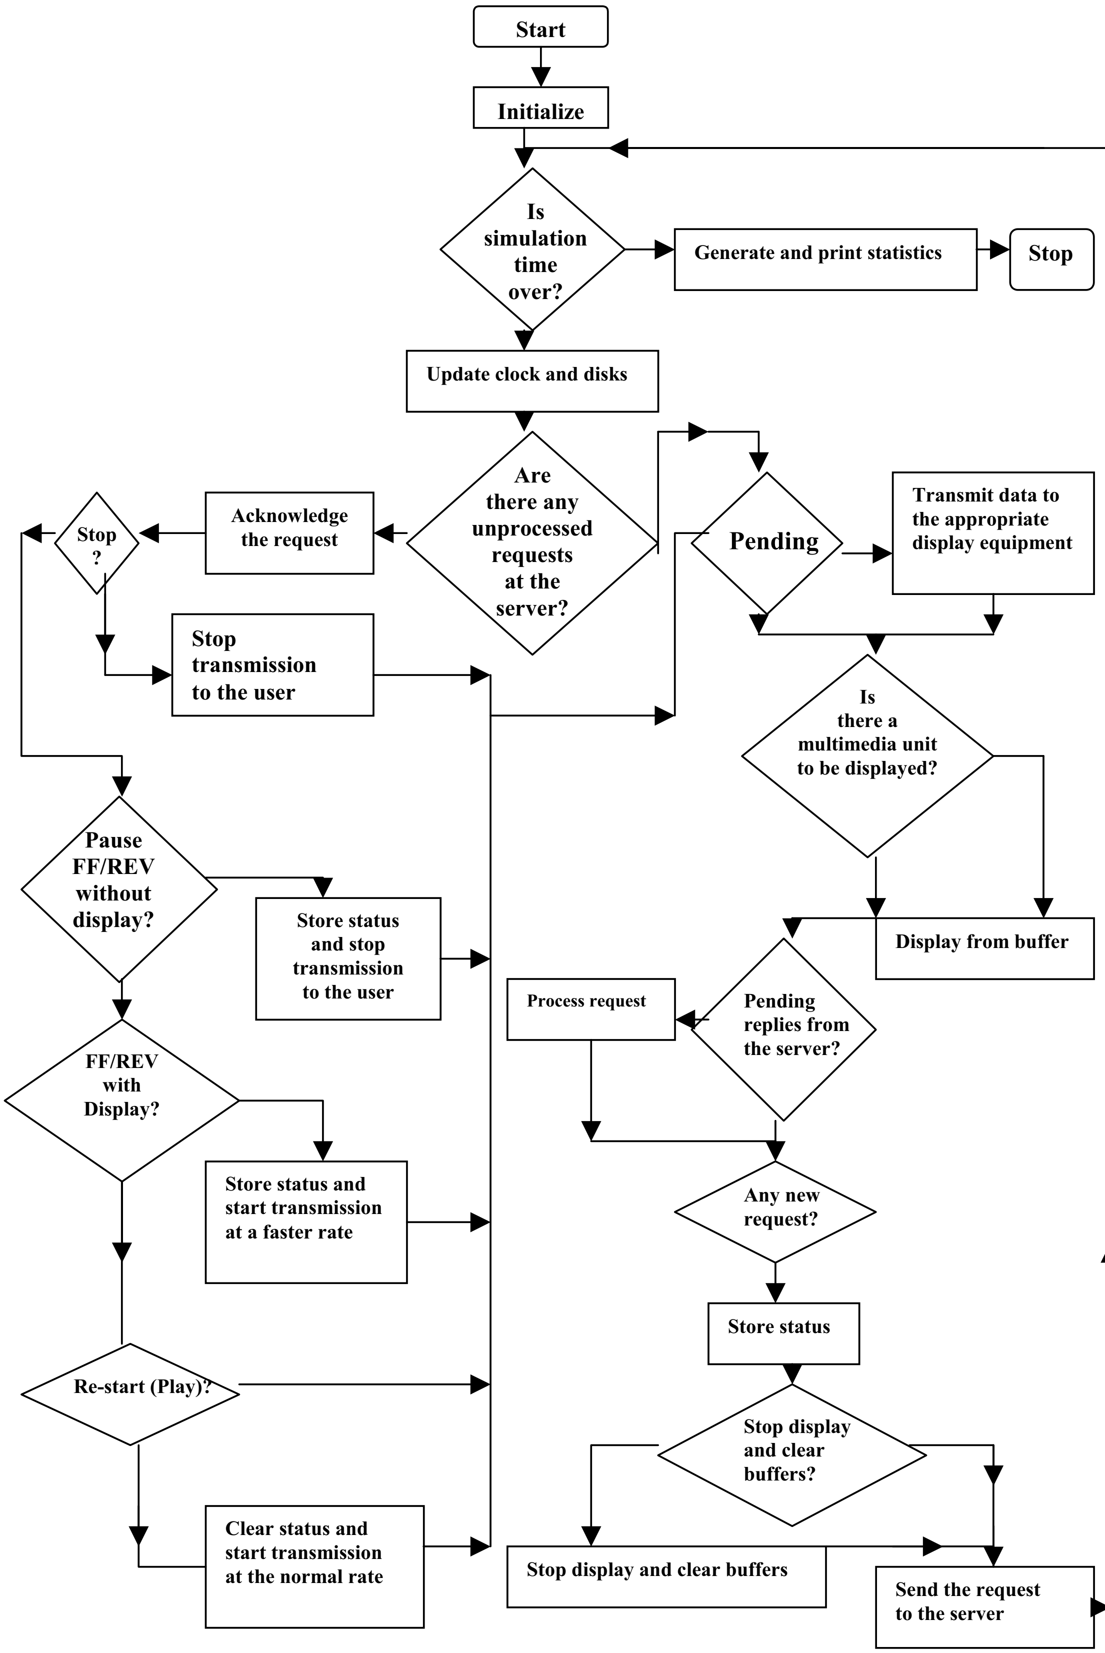

TABLE 1

Comparison of Buffer Requirement in Contiguous and Constrained Allocation Schemes

Network Delay	*Movie Time*				
	1 Min.	*2 Min.*	*3 Min.*	*4 Min.*	*5 Min.*
0.02 Sec	0 2,122	0 4,783	0 7,295	0 9,412	0 12,014
0.04 Sec.	0 4,752	0 9,612	0 14,531	0 18,952	0 23,760
0.06 Sec	0 7,231	0 13,784	0 –	0 –	1 –
0.08 Sec	0 9,412	0 18,983	0 –	0 –	1 –
0.20 Sec	0 –	0 –	0 –	1 –	1 –

each cell, the upper value is the buffer requirement for constrained allocation and the lower value is for contiguous allocation. In *Table 1*, a dash is used in place of the actual number when the number of buffers become greater than 24,000. As we can see from *Table 1*, the number of buffers required for the contiguous allocation scheme is much larger than the buffers required for constrained allocation. Also, the buffer requirement for contiguous allocation increases rapidly as the movie progresses. An increase in network delay also increases the buffer requirement. From *Table 1*, we can see that the buffer requirement for contiguous allocation, with the display of just five minutes of the movie becomes 23,760, whereas for contiguous allocation, the number of buffers required increases only to 1. Since increasing buffers increases the cost of the system, the contiguous allocation scheme is not economically feasible for any full-length movie.

Table 2 shows the variation in buffer requirement for a constrained allocation scheme with an increase in the network delay. A movie of 100 minutes in length is considered. The buffer requirement increases with an increase in the network delay.

6.2 Start-Up Delay

Figure 6 shows the start-up delay for a constrained allocation scheme. The number of users is kept constant at 100, and the number of normal disks used per movie is increased from one to five. The start-up delay is calculated for a 100-minute long movie. It can be observed from this figure that as the number of disks increases, the start-up delay decreases. The decrease in start-up delay is significant when the number of disks is increased from one to two, but it becomes more gradual as the number of disks is increased. The reason for this is that according to our design, the start-up delay is inversely proportional to the number of disks. Thus, when the number of disks is increased from one to two, the start-up delay is effectively halved [5]. By considering a cost curve with this figure, the optimal number of disks required for a desired start-up delay can be found.

6.3 Delay in a Remote Control Feature (Pause/FF/REV) to Become Active

A delay is experienced before a selected remote-control feature (Pause/FF/REV) becomes active because, for some of the remote-control features (features with display), the provision of multimedia units to the display equipment from a special disk is required [5]. This involves locating an appro-

TABLE 2

Buffer Requirements for Constrained Allocation Scheme

Network Delay	*Movie Time 100 Min.*
0.02 Sec	2
0.04 Sec	3
0.06 Sec	3
0.08 Sec	4
0.20 Sec	7
0.25 Sec	9
0.30 Sec	11

FIGURE 6

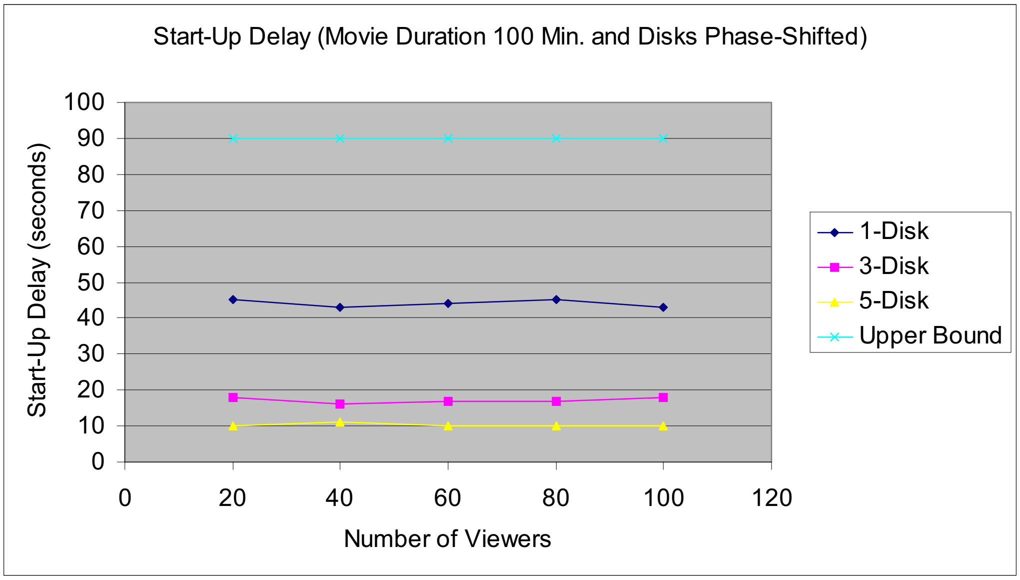

priate special disk and finding the desired multimedia unit. This activation delay is inversely proportional to the number of special disks. *Figure 7* illustrates the activation delay for a constrained layout scheme. The inverse relationship between the activation delay and the number of special disks is also illustrated in this figure by varying the number of special disks from one to five.

7. Conclusions

In this paper, an object-oriented generalized simulation model is developed for multimedia-on-demand systems. This generalized model can be easily extended to simulate various services such as movie-on-demand, interactive network games, remote browsing, multimedia library access,

FIGURE 7

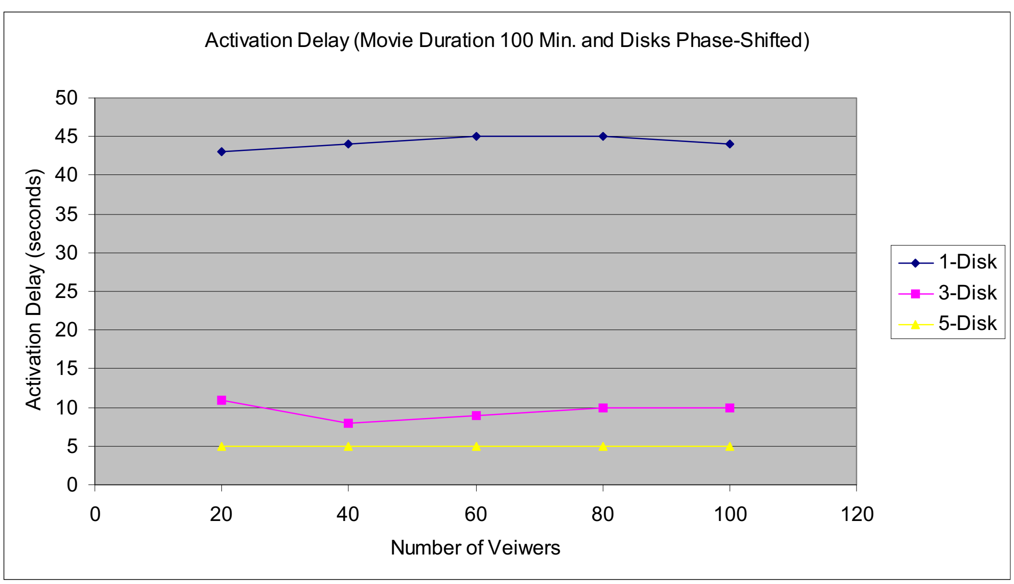

and remote education. Different algorithms for implementing these services can be compared using the simulation model. One of these studies is conducted in this paper, where two disk layout schemes (continuous and constrained) are compared. The comparison is based on buffer requirements for each approach. It can be observed that constrained allocation provides better results than continuous allocation. In addition, the delays for constrained allocation are also reported.

References

[1] Diane P Bischak and Stephen D. Roberts, "Object-oriented Simulation," Proceedings of the 1991 Winter Simulation Conference, 194–203.

[2] Arturo I. Concepcion and Bernard P. Seigler, "DEVS Formalism: A Framework for Hierarchical Model Development," *IEEE Transactions on Software Engineering*, Vol 14. No. 2 (Feburary 1988).

[3] Bernard P. Zeigler, "Object-Oriented Simulation with Hierarchical, Modular Models, Intelligent Agents, and Endomorphic Systems," (Boston: Academic Press Inc. & Harcourt Brace Jovanovich, 1990).

[4] Cheng-Jye Luh and Bernard P. Zeigler, "Abstraction Morphisms for World Modeling in High Autonomy Systems," Proceedings of the 1991 Winter Simulation Conference, 1,129–1,138.

[5] Aditi Singru, Alok Srvastava, and Anup Kumar, "Design of Multimedia-On-Demand Systems," Technical Report AK-EMCS-011-1994.

[6] M. De Prycker, W. Verbeist, and H. Verhille, "Video on Demand over ADSL: Triggering B–ISDN Introduction" (1994).

[7] P. Venkat Rangan, Harrick M. Vin, and Srinivas Ramanathan, "Designing an On-Demand Multimedia Service," *IEEE Communications* (July 1992), 56–64.

[8] Sadahiko Kano, Ken'ichi Kitami, and Masatoshi Kawarasaki, "ISDN Standardization," Proceedings of the IEEE Special Issue on ISDN, Vol. 79, No. 2 (February 1991) 118–124.

[9] M. Gerla and A. Pazos, "Bandwidth Allocation and Routine is ISDNs," *IEEE Communications* (February 1984).

Exploring the Role of Ethernet-Based Technologies in the Metro/Access Network

Ron Landi
Senior Engineer, Network Architecture and Advanced Technology
WorldCom

1. Introduction

This paper explores the role of Ethernet-based technologies in the access and metropolitan networks. There are many changes occurring in the carrier access and metro space making Ethernet-based solutions an appealing and potentially viable alternative to traditional networking technologies. It is important for carriers to understand how access is changing: the transition from narrowband to broadband applications, the increasing move from traditional voice to data and media-centric applications, and the emergence of new technologies focusing on the access and metro environments. There are many new applications and communications services enabled by these changes. Examples are e-commerce, Web hosting, bundled voice and data, mobile data, content/media management, bandwidth-on-demand capabilities, and cell- or packet-based voice applications (e.g., voice over Internet protocol [VoIP]). Many telecom industry forecasts and predictions suggest that Internet (IP) and data traffic will continue to grow significantly over the near- and mid-term periods—with data becoming the dominant type of traffic carried over the carrier networks. Newly emerging legacy-free carriers are already building high-speed optical networks using Internet protocol (IP)–based protocols, while existing carriers are exploring these technologies and beginning to transition their legacy networks. This paper explores the fundamental differences and drivers impacting the network architecture, technology, and costs of Ethernet-based networks as compared to traditional networks.

2. Traditional Verses Ethernet-Based Networks

2.1. Traditional Networks

Today's carriers primarily leverage time division multiplexing (TDM) and asynchronous transfer mode (ATM) technologies to transport voice and data across customer sites in the access and metro networks. The following are some network considerations of these traditional networks:

- They place costs in the International Organization for Standardization (ISO)/Open Systems Interconnection (OSI) network model layers 1 (physical) and 2 (data-link) for service control and network intelligence
- These networks provide carrier-level quality of service (QoS)
- They provide the highest levels of network reliability (i.e., five nines)
- These networks scale as the customer base grows
- These networks have proven network performance, security, and control capabilities

Figure 1 provides a general view of the traditional network. The building houses multiple customers who are served by the carrier. Connecting the building to the carrier network is an access ring—this could be a synchronous optical network (SONET) ring or an asynchronous transfer mode (ATM) data ring. The access ring terminates into the building via access equipment that may be located in the building basement or service-provider point of presence (POP). The access equipment terminates an optical signal, converts to electrical signals, and demultiplexes the channels into digital signal (DS)–3s. The DS–3s are provided, via in building wiring, to each of the customer locations in the building. The DS–3s may be channelized time division multiplexing (TDM) or concatenated ATM transport circuits terminating to customer equipment. The local terminal is an aggregation point in the carrier metro network. Multiple buildings and access rings will terminate into this location. The local terminal aggregates, multiplexes, and switches traffic; in addition, it transits traffic to the long-distance (LD) terminal via wavelengths across the metro/regional ring. The LD terminal provides additional aggregation, multiplexing, and switching functions, in addition to handing off traffic to the long-distance/backbone network.

2.2. Ethernet-Based Networks

During the past one to two years, legacy-free carriers have been entering the access/metro market space. These legacy-

FIGURE 1

Today's Metro TDM/ATM Networks

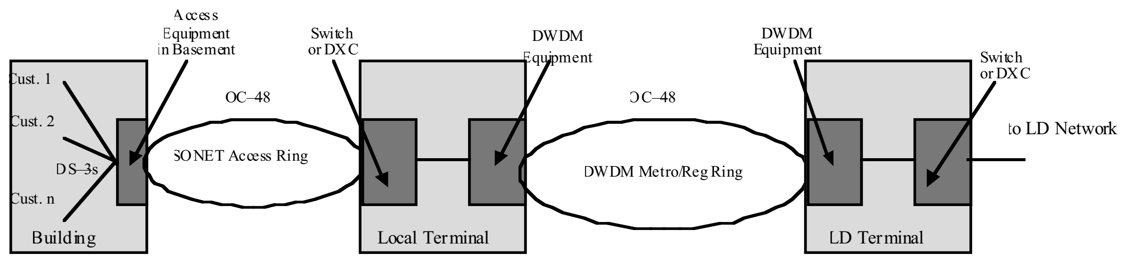

free carriers are deploying Ethernet-based technologies to transport traffic across customer sites in the access and metro networks. These carriers and some large enterprise companies are expanding their local-area networks (LANs) into the access and metropolitan networks. Ethernet is a physical-layer LAN technology and the de facto protocol for IP traffic—speed, cost, and ease of implementation have driven its popularity as a LAN technology in the enterprise market, and it is beginning to take hold in metropolitan-area networks (MANs). The Institute of Electrical and Electronics Engineers (IEEE) LAN standard has been around since 1980, and it is one of the IEEE 802 LAN standards impacting layers 1 and 2 of the ISO/OSI network model. *Figure 2* shows a portion of the ISO/OSI model stack and where Ethernet resides. The following are some network considerations:

- Ethernet takes costs out of the ISO/OSI network model layer 2 (data-link) and relies on layers 3 (network) and above (transport, session) for control and intelligence.
- Shared Ethernet lacks carrier-level QoS and provides "best efforts" and tiered services.
- Ethernet makes "plug and play" of networking equipment a reality but addresses only IP data.
- More than 90 percent of data originates on Ethernet framing.
- Ethernet is expected to grow at approximately the IP growth rate, forecasted at 50 to 100 percent per annum, but Ethernet growth will be limited by access to and availability of fiber in the access and metropolitan networks.
- Ethernet networks have some cost advantages, such as local-exchange carrier (LEC) bypass by moving customers on-net, lower-cost equipment and ports, and potentially lower operating costs to support a less complex network.
- Ethernet has potential for greater efficiencies in networks carrying IP traffic.

Figure 3 provides a general view of the Ethernet-based network. The building houses multiple customers who are served by the carrier. Connecting the building to the carrier network is a dark fiber access ring. The access ring terminates into the building via a Layer-2/3 Ethernet switch that may be located in the building basement or service-provider POP. The Ethernet switch terminates Gigabit Ethernet (GbE) signals transported over fiber. The Layer-2/3 Ethernet switch demultiplexes the GbE signal into 100 Mbps signals, providing 100 Mbps Ethernet connections to the customer routers. The local terminal is an aggregation point in the carrier metro network. Multiple buildings and access rings will terminate into this location. The local terminal aggregates, multiplexes, and switches traffic via another larger Ethernet switch. In addition, it transits traffic to the LD terminal via GbE over fiber across the metro/regional ring. The LD terminal provides additional aggregation, multiplexing, and switching functions in addition to handing off traffic to the long-distance/backbone network.

Ethernet is being deployed today to provide Internet access and point-to-point LAN connections. In the future, it will support multipoint transport between multiple sites. Ethernet alternatives today include the following:

- 10BaseT/100BaseT for enterprise networks. These networks leverage established and widely used LAN tech-

FIGURE 2

L3 Network	**Internet Protocol**
L2 Data Link	**ATM, PPP, *IEEE 802.3 Ethernet***
L1 Physical	**SONET, *Gigabit Ethernet***
L0 Medium	**Fiber, WDM Wavelengths**

FIGURE 3

Ethernet Networks

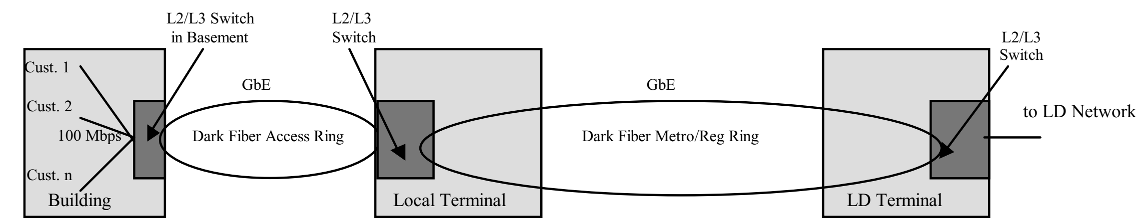

nologies. They offer potentially cheaper point-to-point access and aggregation in the access/metro network.

- Gigabit Ethernet: Ethernet carried over optics, such as SONET or dark fiber rings. These networks are designed for interoffice links and high-speed LANs, extending the LAN into the MAN.

Ethernet alternatives in the future include the following:

- 10 Gigabit Ethernet: Ethernet over SONET Lite. 10 Gigabit Ethernet is an evolving solution for metro and core (or backbone) transport. The standards are being developed at this time, with the first switch ports due out by the vendors in 2002. This Ethernet solution should provide scalable IP over Ethernet over SONET transport, enable improvements in transport of IP data traffic, and enable improved efficiencies in the regional and core networks through flow-control capabilities.

3. Ethernet/GbE Market Size and Potential

The U.S. market for Ethernet-based services is potentially very large. The total U.S. interconnect market size in 2000 was estimated at about $38 billion—this includes all LAN private lines, frame relay, and ATM. The total U.S. interconnect market size is forecasted to grow to more than $100 billion by 2005. In 2005, a reasonable portion of this interconnect will be via Ethernet-based transport—the market addressable via Ethernet is estimated at more than $20 billion in 2005. The build-out and availability of fiber in the access and metro networks is a key factor limiting this estimate. Gigabit Ethernet services, including Internet service provider (ISP) access and the U.S. interconnect market, is estimated at more than $2 billion in 2000 and is projected to grow to $10 billion by 2005. Packet solutions, such as Ethernet, will enable growth in the access and metropolitan networks. As depicted in *Figure 4*, in 2000, circuit-centric technologies dominate the access network and packet-centric technologies dominate the edge/core networks. In 2005, as applications and new services drive up the demand for bandwidth across the network, it is anticipated that packet-centric technologies will be dominant across the network—in the access, metro, edge, and core networks.

4. Players in the Ethernet/GbE Space

There are a number of emerging service providers that are focusing on Ethernet-based networking, specifically in access and metro city networks. Existing network infrastructure and voice-centric equipment do not encumber these legacy-free carriers. They are leasing facilities and dark fiber, deploying Ethernet equipment, and targeting

FIGURE 4

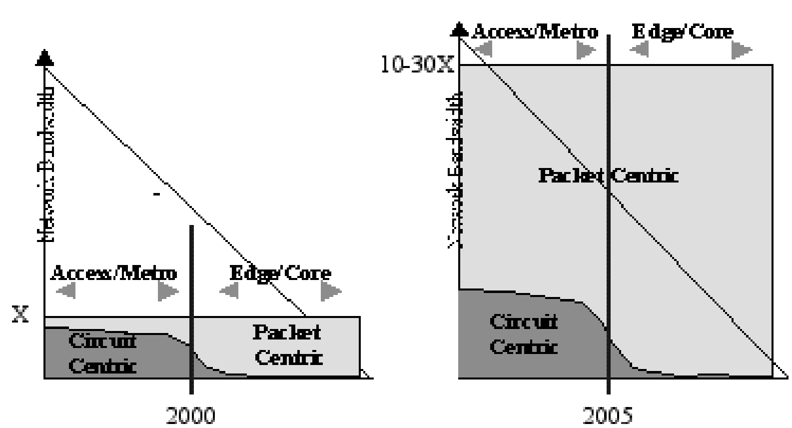

specific local markets offering data services. These new legacy-free carriers have the following characteristics:

- They are small networks serving hundreds of customers.
- Most are leasing dark fibers in MANs.
- Most are outsourcing the long-haul transport via traditional (SONET) solutions.
- They are offering data services, such as high-speed Internet access, point-to-point dedicated, and best-effort shared services.
- Gigabit Ethernet does not scale presently—limiting metro network capacity and making monitoring/supporting these new data networks difficult for the legacy free carriers.
- Significant growth will be enabled by 10 Gigabit Ethernet and network (ISO/OSI model layer 3) improvements. This will extend the carrier Ethernet-based network to the regional and core networks—in the future.

Figures 5a through *5c* depict the basic network deployments of these new legacy-free carriers. *Figure 5a* shows a topology that supports a shared Ethernet service. This network interconnects Ethernet Layer-2/3 switches that provide 100BaseT Ethernet interfaces to the customer routers. These switches can do traffic policing, rate limiting, and shaping, allowing the service provider to provide a managed service with varying bandwidth. The customer is provisioned a physical port interface and could be sold committed and burstable bandwidth. The Ethernet frames are aggregated at the edge and transported as GbE over dark fiber or optical (e.g., dense wavelength division multiplexing [DWDM]) rings to a router that connects to private and public Internet networks. *Figure 5b* also shows a shared Ethernet topology that maps Ethernet frames into SONET/TDM payloads. This network interconnects SONET add/drop multiplexers (ADMs) that provide 100BaseT Ethernet interfaces to the customer routers. The service provider provides a managed service with a fixed amount of bandwidth delivered to the customer. The customer is provisioned a physical port interface and sold committed bandwidth. The Ethernet frames are aggregated at the edge, mapped to SONET payloads, and transported across the metro network over a SONET ring to a router that connects to private and public Internet networks. *Figure 5c* shows a dedicated Ethernet service using a meshed network topology. This network interconnects Ethernet Layer-2/3 switches that provide 100BaseT Ethernet interfaces to the customer routers. These switches can do traffic policing, rate limiting, and shaping, allowing the service provider to provide a managed service with varying bandwidth. The Ethernet frames are aggregated at the edge and transported as GbE over dark fiber or optical (e.g., DWDM) connections to routers that connects to private and public Internet networks. This network alternative provides an additional physical path, from the customer site to the carrier network, for redundancy and restoration requirements.

5. *Ethernet Networking Options*

The basic Ethernet networking options presently being explored by carriers include the following:

- Ethernet private lines supporting IP over Ethernet over SONET (refer to *Figure 6*). This network configuration provides dedicated capacity and private line (Ethernet frames mapped to SONET payloads) services.
- Shared Ethernet networking supporting IP over Ethernet over SONET (refer to *Figure 7*). This network configuration provides shared capacity, shared

FIGURE 5

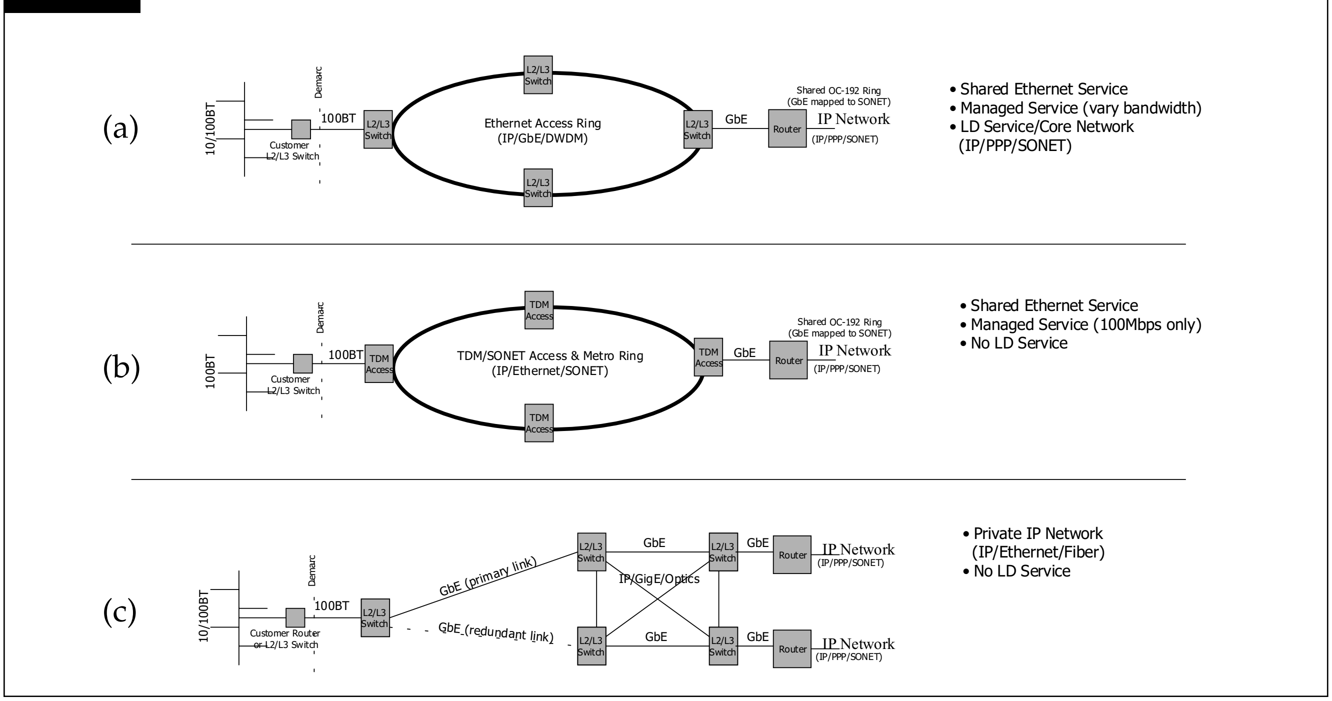

Ethernet over SONET, with full bandwidth available to all customers (i.e., best-effort service).

- Shared Ethernet networking over fiber supporting IP over Ethernet over fiber (refer to *Figure 8*). This network configuration provides shared capacity and full bandwidth available to all customers (i.e., best-effort service). It also introduces potential network operational issues involving the scaling/growing of customers on network and fiber management.

These networks support the following services: high-speed Internet access (IP), shared and dedicated point-to-point services, and shared and dedicated multipoint services. The point-to-point and multipoint services support private corporate/enterprise LANs and metro access networks, with secured interconnections to external and public networks.

Figure 6 shows a network topology supporting dedicated Ethernet services. This network leverages existing access and metro network rings, digital cross-connect (DXC)ing equipment, and metro DWDM transport equipment. A new access platform or native Ethernet interfaces on existing SONET ADM platforms are deployed at the customer location or POP site. These platforms terminate the fiber rings into the POP site, perform SONET payload and Ethernet frame mapping functions, and may perform limited multiplexing functions. In this example, multiple customers are provided 100 Mbps dedicated connections to routers or small Ethernet switches. The customer traffic is mapped to specific TDM time slots and transported across the SONET network—the customer traffic is segregated from other traffic and maintains private-line characteristics. The incentive for this type of deployment is a less-expensive private-line alternative service that is easy to implement and requires minimal operational systems modifications. *Figure 7* shows a network topology supporting shared Ethernet services. This network transports Ethernet transported as point-to-point protocol (PPP) over SONET. This network uses Layer-2/3 Ethernet switches at the POP sites to provide Ethernet traffic policing, rate limiting, and shaping capabilities. The customers are provided 100 Mbps Ethernet interfaces that connect to their routers or smaller Ethernet switches. They may be sold a committed amount of bandwidth with a burstable level—the level of bursting will vary based upon the customer's actual usage, contention with other customers in a given time period, oversubscription factors applied to the noncommitted bandwidth, and provisioned trunked capacity. This network relies on larger Layer-2/3 switches in the carrier network to effectively aggregate, groom, and multiplex traffic for more efficient transport. *Figure 8* also depicts a network topology supporting shared Ethernet services. This network transports Ethernet transported as PPP over dark fiber. This alternative, while potentially cheaper than SONET transport, may yield operational issues regarding scaling (growing/expanding) the network and management of the physical fiber infrastructure.

6. Cost Comparisons of Networking Alternatives

This section and the next section discuss the Ethernet networking costs as compared to traditional data-networking alternatives. A baseline model, IP over TDM over SONET, serves as a comparator to the other network options. The building or POP houses the customers. The access ring, between the building and local terminal, drops an OC–48 of capacity to the building where it is terminated via SONET ADM and associated equipment. The access ring is one fiber pair using 1310 nm wideband optics. The OC–48 ports on the broadband DXC (BBDXC), and SONET ADM are protected. At the building ADM, the OC–48 is converted to electrical signals and demultiplexed to a DS–3 level. Each of the 48 customers in the building receives a DS–3 to his or her customer router equipment. The customer router converts the DS–3 and distributes up to twenty 10BaseT and 100BaseT interfaces to the customer personal computers (PCs) and other Ethernet-based appliances. All OC–48 and DS–3 interfaces are TDM. This baseline model assumes only data services are delivered to the customer. *Figure 9* depicts the baseline network model.

FIGURE 6

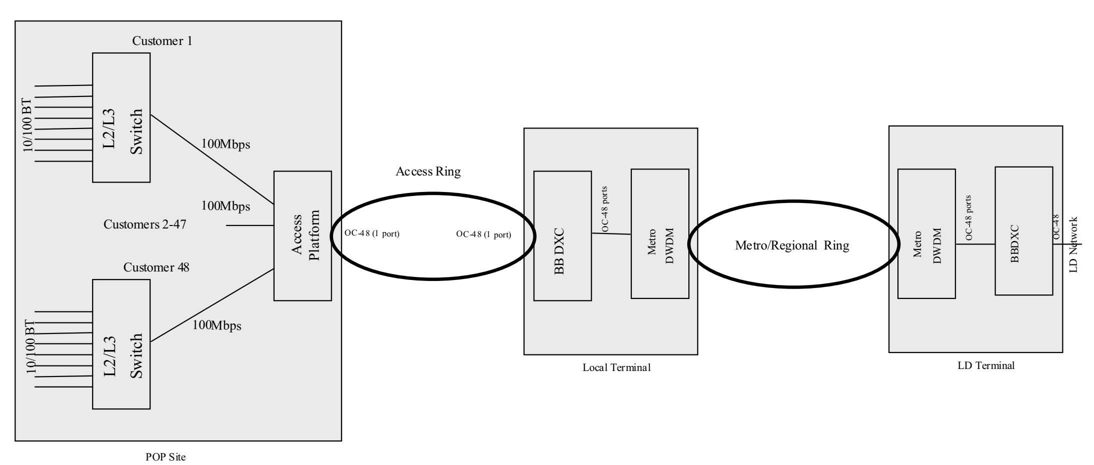

FIGURE 7

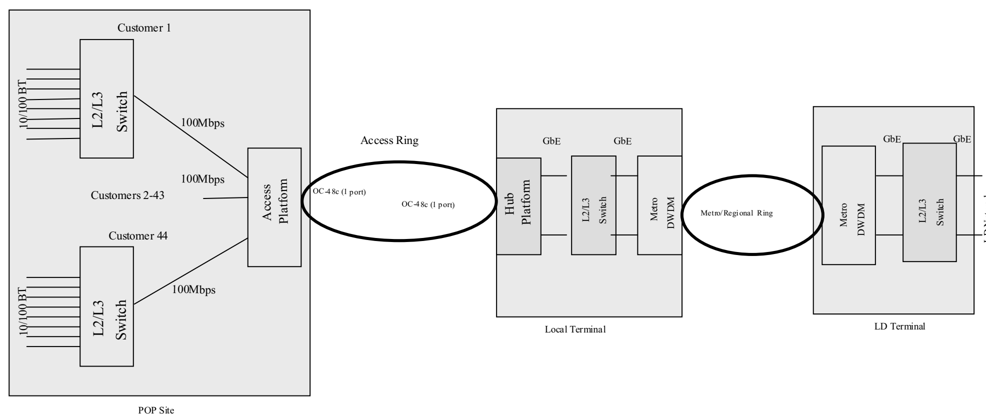

FIGURE 8

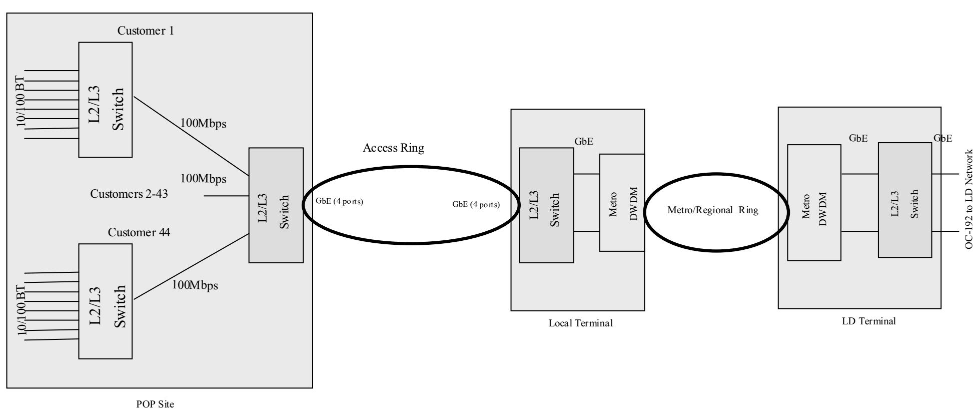

7. *Ethernet/GbE Considerations and Economic Drivers*

A high-level comparison of network capital costs for transporting IP traffic across various networks is included in *Figure 10*. The IP/TDM/SONET model is the baseline model used by many of the legacy carriers today. This model transports IP (layer 3) over PPP (layer 2) over SONET (layer 1) transport—the SONET could be carried over fiber or WDM (layer 0). This model provides TDM or other channelized interfaces to the customer. This model is compared with four other network models. The first model is IP/ATM/SONET, with ATM as the layer-2 technology. ATM is a connect-oriented protocol that provides QoS, flexibility, and supports voice, data, and video services. This model is approximately 46 percent higher in network capital costs than the baseline model (IP/TDM/SONET)—higher costs result from network equipment and interfaces. The second model is Ethernet private line. This model maps Ethernet frames to SONET payloads, with the customer traffic assigned to specific time division slots. The model is similar to the baseline model, except that the customer is provided with a native Ethernet interface (e.g., 10BaseT, 10/100BaseT, GbE). This model is approximately 32 percent less than the baseline model—lower costs result from cheaper customer interfaces. The third model is shared Ethernet over SONET. This model uses Ethernet Layer-2/3 switches to provide switched services. These platforms also provide policing, rate limiting, and shaping of traffic at the POP site. Customers are allowed to burst beyond a committed information rate and share the trunk/network capacity. This model is approximately 41 percent cheaper than the baseline model—lower costs result from lower network and customer equipment costs. The fourth model is shared Ethernet over dark fiber. This model is similar to the third model, except that the Ethernet traffic is transported directly over dark fiber, without Layer-1 SONET. This model is approximately 49 percent cheaper than the baseline model. Some key assumptions of the models include that the carrier builds fiber plant (access and metro/regional rings); no statistical multiplexing gain is applied to cell- and packet-based alternatives; the carrier uses WDM in metro/regional transport rings; and equipment list prices and most recent equipment types are used.

Depending on the volume, nature (voice, data), type of services supported, and customer requirements (e.g., service-level agreements [SLAs]), it is possible to lower the costs in the IP/ATM/SONET and shared Ethernet models. This

FIGURE 9

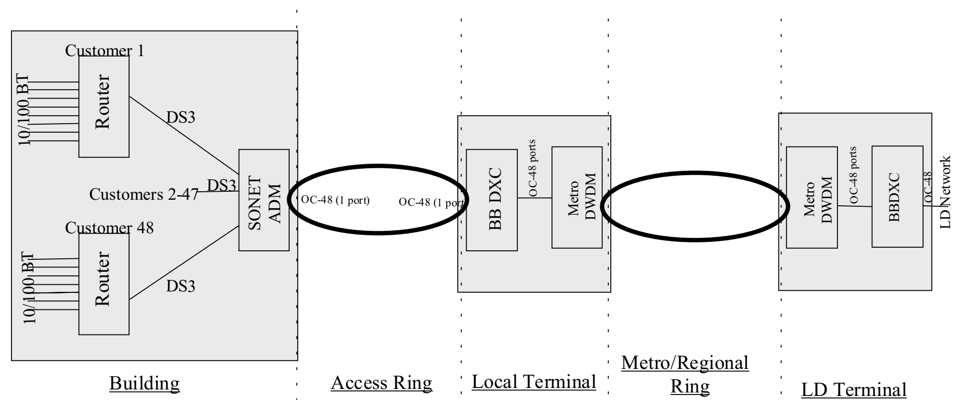

could be realized by applying statistical multiplexing to the shared trunk/network capacity. An over-subscription of 2:1 or 4:1, depending upon the services offered and the customer's requirements, would reduce the network costs. There would be transport savings in addition to DXC and switch port savings.

Some of the reasons for cost differences between the models include the following:

- Shared lines translate to shared costs and less idle bandwidth compared with dedicated transport services.
- TDM private lines
- Ethernet network is seamless and requires no protocol conversion—translating into simpler and less equipment.
- Ethernet equipment is designed for "best-effort" services and is less costly than telco (carrier class) reliable equipment.
- Potentially lower operational costs, dependent upon deployment and legacy infrastructure, may be realized by the carrier—such as the following:
  - WAN verses LAN (cheaper) technical labor and support
  - Faster provisioning of additional bandwidth to existing customers
  - Simpler network monitoring and performance
  - Integrated and simpler back-office systems
- Incremental pricing is possible with Ethernet and Gigabit Ethernet transport services, resulting in finer granularity of bandwidth to be sold to customers by the service provider. Examples include the following:

FIGURE 10

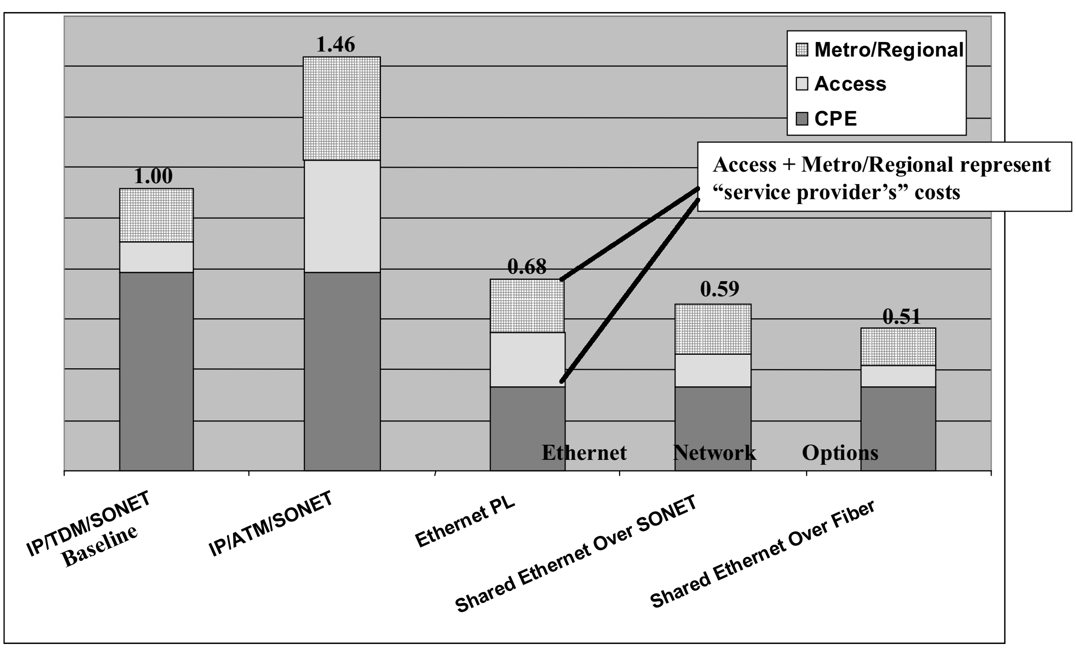

- 10 Mbps Ethernet interface (bandwidth sold in 1 Mbps increments)
- 100 Mbps (fast) Ethernet (bandwidth sold in 5 Mbps increments)

8. Technology Comparisons of Networking Alternatives

In comparing traditional networks, such as TDM and ATM, to Ethernet-based networks, it is important to consider a number of factors. The following outlines some of the basic trade-offs of the different technology choices:

- *Pros of Traditional (TDM, ATM) Networks*:
  - Carrier-level QoS
  - Five-nines reliability
  - Technology platform scales as network (customer base) grows
  - SONET performance, security, and control
  - Supports legacy (high-margin) services
- *Cons of Traditional (TDM, ATM) Networks*:
  - High-cost customer equipment
  - Complex network management
  - Difficult and time-consuming to provision incremental bandwidth and new services
  - Not IP (data) efficient
  - Difficulty in supporting forecasted high growth of IP traffic
- *Pros of Ethernet-Based Networks*:
  - Cheap customer ports/interfaces
  - Reduced customer-premises equipment (CPE)
  - Potentially lower operational costs
  - Increased bandwidth control and efficiencies
  - Supports faster provisioning of new services and bandwidth
- *Cons of Ethernet-Based Networks*:
  - Provides enterprise, not carrier-level, QoS
  - Provides three-nines reliability
  - There are no line-, path-, and section-monitoring capabilities
  - Supports IP only services
  - Requires fiber access
  - Ethernet-based standards are evolving
  - There is an embedded base of legacy/TDM network infrastructure

Some advantages of Ethernet-based networks over SONET and ATM networks include the following:

- *Flexibility:* Ethernet can be purchased in small increments rather than large chunks of bandwidth
- *Provisioning:* Bandwidth provisioning can be done in hours, not days or weeks. Web-enabled capabilities allow the customer to change bandwidth as application needs change
- *Cost:* Less equipment and lower operational costs mean that Ethernet-based networks can be priced at less than SONET and TDM
- *Efficiency for Data:* Ethernet was designed for data transmission, requiring no protocol conversion for traffic that is increasingly data-centric

Figure 11 illustrates the inter-relations between technology and business issues regarding Ethernet-based technologies and networks. There are a number of technology issues, including standards, costs, network simplicity, network efficiencies, scalability, and reliability. There are a number of business issues to be considered. These issues include markets to be served by Ethernet-based technologies, services supported (e.g., voice, data, video), competitive factors relating to newly emerging legacy-free carriers and existing service providers, service-provider partnering and joint-

FIGURE 11

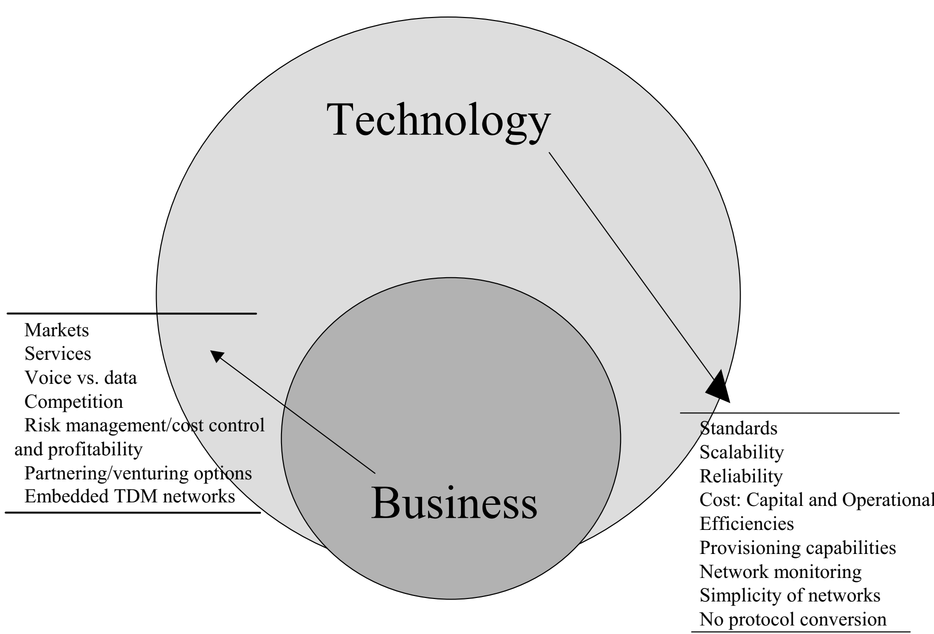

venture opportunities, and the legacy carriers' existing TDM/ATM–based networks—how is Ethernet migrated into the existing network? The carrier must carefully consider and weigh the technology and business issues in making the decision to proceed with Ethernet-based technologies and networks. For the legacy carrier, this decision can be more complex—balancing existing infrastructure investments and customers against new investments, capabilities, and services.

9. Summary and Conclusions

Much of the industry and market hype during the past year or so regarding Ethernet-based network costs is false. Ethernet-based networks are not 1/10 to 1/15 the cost of traditional networks. The capital cost for Ethernet-based networks is approximately 50 to 60 percent less than traditional networks. Other realities of Ethernet-based networks include the following: fiber access makes Ethernet-based network deployments capital intensive; Ethernet is IP–centric and not integrated with legacy services; Ethernet networks don't scale today; and Ethernet QoS levels need improvement to be viable for large national and global carriers. Ethernet-based technologies will grow in number of deployments, breadth of network (moving from the LAN to the MAN to the WAN), and number of customers served. This growth is already happening and will continue to be fueled by the following developments:

- Growth of Ethernet will be proportional to growth of IP traffic and related applications.
- Operations support systems (OSS) and service provisioning/delivery will likely be simpler, yielding lower operating costs.
- Cheaper delivered bandwidth cost (Ethernet is 1/3 to 1/2 the cost of traditional networks) and lower operating costs could make this cost difference/savings even more substantial.
- IEEE standards are evolving that will support large networks and their need to scale:
  - 10 Ggabit Ethernet
  - Resilient packet ring (RPR)
  - Standardized methods to map Ethernet frames to SONET payloads (e.g., X.86, generic-framing protocol)

Ethernet-based technologies are making their debut in the access and metro networks today. As this technology alternative matures and its standards evolve, Ethernet-based technologies should begin to take center stage in the near- to mid-term future.

The MxU Market: Reaching over Current Networks Limitations

Haim Laufer
Vice President, Marketing
Infineon Technologies

Introduction

The demand for added-value broadband services to the home and office is still on the rise despite the recent setback in the telecommunications market and world economy. Increasing Internet penetration into daily activities and the growing need for bandwidth in homes and businesses is prompting service providers to implement broadband access solutions.

Driven by competition for the service and application dollar, incumbent local-exchange carriers (ILECs), competitive local-exchange carriers (CLECs), cable and satellite companies, interexchange carriers (IXCs), and Internet service providers (ISPs) are searching for new technologies that are able to deliver very high bandwidths in a cost-efficient manner. Targeting clustered environments, such as the multitenant/multidwelling unit (MxU) market, and especially those that can be approached as a group (hotels or office buildings), offer attractive business opportunities where the penetration costs required are justified by the potential revenues from services. Broadband access options to the building include fiber optics, broadband wireless (local multipoint distribution system [LMDS]), cable television, or satellite. From this point of building entry, broadband services must be distributed inside the building. Since the basic services to be offered are Ethernet-based, an Ethernet network inside the building becomes the most natural and cost-effective solution.

The MxU Market

A large potential market for MxU installations is in office buildings. Since businesses tend to send as much data as they receive, they need symmetric broadband access for the office local-area network (LAN) environment as well as to connect to the outside world. Residential users, on the other hand, demand entertainment services such as gaming and video. These types of services require very high bandwidths, mostly asymmetric, although services requiring a more symmetric solution—such as music or video-files sharing, photo albums, home networking, etc.—will also make their entry into the residential MxU sector in the future.

Another rewarding MxU segment is the hospitality market. Hotels and lodging facilities have to cost-effectively supply what today's business traveler needs. This goes beyond "bed and breakfast" and extends all the way to the Internet and global communications. Although the wiring infrastructure is already in place, it is not designed to support a LAN to the guest rooms. Therefore, the proprietors must use a solution based on LAN technology with as little trade-off for bandwidth and functionality as possible with the fastest installation, yet avoiding the costly change of their wiring, and with minimum disruption to the normal hotel operation.

Figure 1 shows the anticipated revenues from MxU hardware and services.

Ethernet In-Building Network Solution

Deploying a standard Ethernet network inside the building seems straightforward. Independently of the transport media used to bring the broadband pipeline to the building—fiber, broadband wireless, or satellite—the inside network will consist of a similar architecture: a 10/100BaseT port provided for each unit, whether a hotel room, apartment, or office. Due to range limitations, a hub must be installed at each floor or every two floors, and all of them should be connected to a switch installed (most probably) at the building basement. The Ethernet switch in the basement serves as an aggregator, connecting the ports to the transport line, at a few 10s of Mbps up to 1 Gbps, depending on the concentration ratio and the media used. When the services/users justify it, an in-building server may be used to reduce traffic and increase the efficiency or to provide local added-value services.

This architecture explained has its well-known advantages: It's simple, easy to install, and uncomplicated to manage. But it also carries some major disadvantages as well. In most of the cases, a Cat 5 cable infrastructure is unavailable and needs to be deployed, especially in hotels and residential buildings. In other cases, it is impossible to install new wiring due to lack of ducts or opposition from the building owners. And in others still, it is very difficult and plain costly. Installing hubs or small switches at each floor may

FIGURE 1

Anticipated Worldwide MxU Hardware and Services Revenues

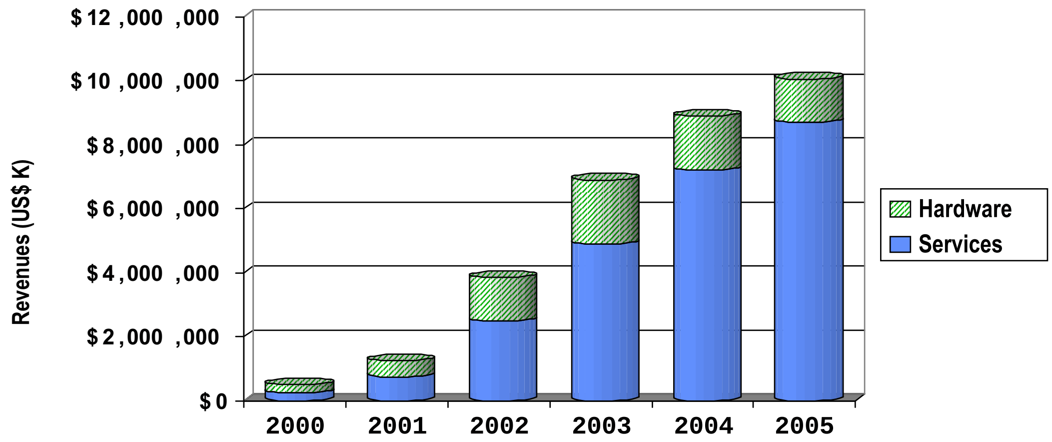

Source: Cahners In-Stat Group

also be a very difficult and frequently impossible task. A locked room must be prepared for security reasons; a power plug and air conditioning are required—and these are only some the problems encountered in the set-up stage.

A Novel Proposal for MxU Network Applications

There is another possibility to enjoy the advantages of simplicity, manageability, and cost-effectiveness of the Ethernet network solution while avoiding the aforementioned disadvantages. The basic idea is to use a technology that can be implemented over the standard telephony twisted-pair available in all buildings, offering symmetric 10 Mbps over a distance of several hundreds of meters. This technology has a name—it's called 10BaseS™.

Based on very-high–data rate digital subscriber line (VDSL), Infineon Technologies' 10BaseS™ solution merges the Ethernet signal on top of the telephony signal using the latest coding and digital modulation techniques, interfacing smoothly and seamlessly with existing Ethernet solutions. 10BaseS systems take advantage of the unlimited potential of existing telephone wire infrastructure to deliver high-speed video, data, and voice services over greater distances. Combining standard Ethernet technology with the robust VDSL technology, 10BaseS™ presents the ideal solution for in-building data distribution applications.

A 10BaseS network is essentially an Ethernet network with an extended range on each 10BaseS segment. The 10BaseS ports of the 10BaseS switches in the basement of the building are connected through splitters to the telephone infrastructure. Nearly one mile away, at the end of each telephone line, another splitter in the 10BaseS customer-premises equipment (CPE) is installed to separate plain old telephone service (POTS), integrated services digital network (ISDN), or private branch exchange (PBX) digital set signals from 10BaseS signals. A standard RJ-45 jack connects one port of the modem to the user's personal computer (PC) or LAN segment. A standard RJ-11 jack connects another port on the modem to the telephone. Appropriate software on central servers provides operation and management functions that enable Internet access and handles billing. *Figure 2* shows several deployment scenarios utilizing 10BaseS technology.

Cost Effective and Revenue Generating

The advantage of 10BaseS is that it uses the existing telephone wire infrastructure. There is no need to rewire the facility with Ethernet grade cables, saving the cost of rewiring and the possible loss of revenue while areas are shut down for laying wires. All the equipment can be installed in the same place, locked, and protected, making the maintenance (including the installation of an uninterruptible power supply (UPS), air conditioning, etc.) easy.

Using existing telephone infrastructures, 10BaseS brings full-speed 10BaseT Ethernet to every telephone extension within a cable mile from the switch, in any facility, without interfering with any existing telephone functionality, based either in regular phone, ISDN, or digital PBX phone. Internet access can even be offered in telephone booths, executive lounges, and cafeterias. Internet access, virtual LAN connectivity with home offices, videoconferencing, locally hosted Web sites, video on demand (VOD), and other data services are all made possible once this Ethernet connectivity is available.

In addition, the residential market offers enormous business opportunities for access service providers, building owners, and ISPs. Today's residential users demand more than just raw Internet and are willing to pay for added-value entertainment services. However entertainment applications—such as interactive games, CD–quality music, and DVD–quality movies on a pay-per-use basis—require bandwidths in the range of 5 to 15 Mbps. Most of today's access systems cannot support this kind of bandwidth or are too expensive to deploy.

FIGURE 2

Broadband Access Applications Utilizing 10BaseS

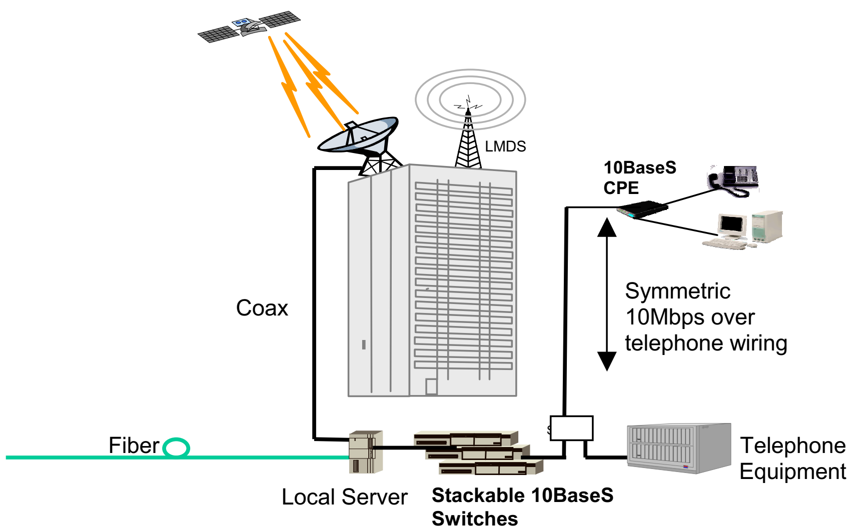

10BaseS can provide each apartment in a residential building, townhouse, garden-style complex, or dormitory with symmetric 10 Mbps network connectivity. Furthermore, taking advantage of VDSL technology, a 26/3 asymmetric scheme can be achieved—more than enough for five video channels per apartment.

Other 10BaseS applications may include public centers such as airports, convention centers, or shopping malls—all generating a lot of traffic. Public-center proprietors and operators utilizing 10BaseS systems will be able to offer high-speed Internet access at every phone booth on the premises. Information kiosks and travel information can be offered through a locally hosted Web server. Offices located on premises can utilize 10BaseS for LAN applications, and security equipment, including closed-circuit surveillance cameras, can be deployed and controlled through a 10BaseS switch connected to the control station—all of this while maintaining legacy voice services unaffected.

A 10BaseS switch located in the building basement or information technology (IT) room makes every phone jack in the building Ethernet-enabled. Coexisting with telephone communications, 10BaseS offers a true broadband connection to the home enabling simultaneous video, voice, and data services over hundreds of meters of telco-grade copper wires.

Conclusion

The MxU market offers service providers a wide range of new business opportunities for distributing video, data, and voice to a large captured customer base. There are a variety of technologies used in the MxU deployments for the in-building applications—Ethernet, without a doubt, being the most popular and simple of them all. Yet Ethernet requires rewiring entire buildings with Cat 5 cables and has a distance limitation requiring the distribution of hubs throughout the building. The only way to overcome Ethernet's disadvantages and still enjoy its simplicity, cost-effectiveness, and bandwidth capacity is by combining it with a technology capable of carrying very high bit rates over greater distances. The 10BaseS™ solution, an innovative technology that offers the simplicity of Ethernet with the high performance of VDSL, does exactly that. A system built with 10BaseS technology delivers full duplex standard Ethernet and telephony service sharing the existing copper wire infrastructure, across a distance of 4,000ft (1,200m) and more. High-speed, noise-free video and data transmission running simultaneously with POTS, ISDN, or PBX signaling are made available with 10BaseS. Taking advantage of 10BaseS, service providers acquire the right set of tools for upgrading their service offerings, achieving better market penetration, and gaining the competitive edge.

How Universal Plug and Play Speeds DSL Deployment and Lowers Service-Provider Costs

Andrew Liu
Business-Development Manager and Co-Chair, Universal Plug and Play Marketing Committee
Intel Corporation

Marc Abrams
Product Manager, Residential Gateways, Residential Access Division
Intel Corporation

This paper presents recent developments in auto-configuring broadband-connected home networks. It recommends that the customer-premises equipment (CPE) vendor shoulder the responsibility for making the end-user experience invisible and automatic to speed deployment and lower costs.

Customers' Requirements

CPE vendors have two customers: digital subscriber line (DSL) service providers and end users.[1] Although DSL service providers are the center of attention for most CPE vendors, the needs of the end user must also be met for successful digital subscriber line (DSL) deployments. While DSL service providers need compatible, cost-effective CPE, physical compatibility is just the first step. The next is flexible and extensible platforms that can be customized for different usage models and can work with complex back-end infrastructures. Going to this next step can result in lower deployment costs for DSL service providers.

End users need equipment that requires no configuration. Asking them to learn to enter Internet protocol (IP) addresses and the like is not acceptable. The work done in the DSL Forum for CPE auto-configuration, such as technical requirement (TR) 37 interim local management interface (ILMI) along with dynamic host configuration protocol (DHCP) is a good first step, but not the endgame. ILMI and DHCP are only enough to provide basic Internet access, which is becoming a commodity business with low margins and high churn rates.

What's needed is a mechanism that can automatically configure more sophisticated services like virtual private networking (VPN), voice over Internet protocol (VoIP), videoconferencing, video on demand (VoD), and the like.

More importantly, there needs to be a mechanism to automatically configure the home networking adapters and personal computers (PC)—the digital home infrastructure—to support advanced services like VoIP and VPN within the home for end users. This is what universal plug and play (UPnP) can do for service providers.

As the kind and number of home-networked devices is changing and growing very rapidly, moving from just PCs to all kinds of Internet-connected devices such as phones, web tablets, and—yes—even the toaster, refrigerator, and microwave oven, a modular hardware and software architecture is advised for CPE to extend the duty cycle and reduce equipment and support costs for the service provider.

Modular Platforms

What does it mean to have such a flexible, extensible, and customizable platform? Offering a low-cost, fixed-function product can be deceptively expensive. Consider the never-ending requests for quotations (RFQ), the CPE qualification process, and even just offering a single CPE to your customers. They all tie up resources and limit the ability to differentiate services and generate new revenue opportunities. Offering multiple CPE drives vendor costs higher and increases training expenses because everyone needs to learn about more equipment. Multiple CPE also challenges the sales and logistics personnel and increases inventory liability for equipment that may not be used.

A modular platform consisting of both software and hardware can solve these problems. When adding a new feature to CPE, such as 802.11b wireless networking, one does not need to requalify the whole platform. The new feature is an add-on. And adding new software features—for example, VoIP functionality—does not require that the modem be requalified.

Having a modular platform can greatly minimize the costs of sales and support training. Service-provider personnel need only learn about the new feature, not a whole new product. And the logistics of inventory becomes less complicated because the modular platform can be used for any kind of installation that requires networking and related services.

DSL Auto-Configuration Is Not Enough
There is a very active special interest group within the DSL Forum working on CPE auto-configuration. Building on the work of TR37 (ILMI) are several new proposals for configuring more and more complex services. These include WT59, WT60, WT65, and WT77. Each builds on the next to help service providers auto-configure even the most complicated types of services and lower operations and support costs.

However valid, the DSL Forum efforts focus on service-provider issues with services configuration. According to Parks Associates Research, an end user is nine times as likely to have a home network as a non-broadband customer. What that means is that once end users get broadband, they want home networking. So how can a CPE vendor meet the no-configuration requirement for these customers if they aren't paying attention? Are vendors addressing the need at all?

CPE Pre-Configuration Is Not Enough
Currently, CPE vendors pre-configure CPE per service-provider requirements. In the cases where an end user needs to provide information to get connected to the network, CPE vendors provide wizard-driven graphical user interfaces (GUI) that direct end users to enter their IP addresses, then enter this, and then enter that.

How does a wizard-driven setup work for a PC? The modem ships with an installation CD for an end user's PC. The installation screen asks if the end user has automatic addressing (DHCP) or is an advanced user who needs some special IP addressing. By answering the questions on a single screen and pressing a few keys, the end user can complete the installation.

Many providers offer a simple setup procedure for DSL modems in which the end user goes to a Web page that the modem hosts. From there, end users take only a few steps and provide a few pieces of information. The steps and questions vary from product to product, but generally at the Start page end users click Begin and answer questions about the type of connection they have. Depending on their service providers, they may have point-to-point protocol over asynchronous transfer mode (PPPoA) or point-to-point protocol over Ethernet (PPPoE) and can have others for more complex services. Also, depending on their service they may or may not have point-to-point protocol (PPP) authentication.

The next screen, which concerns IP, may already be completed if the vendor knows the user and has preconfigured the firmware on the modem. Users who later need to change this information can do so on this screen. The final screen summarizes the modem setup—the results of completing the preceding screens. That is all end users need to do to setup their modems.

Built-in diagnostics let the end user test the modem, the DSL connection back to the digital subscriber line access multiplexer (DSLAM), and the connection to the Internet service provider (ISP). Suppose the ISP fails its test? What the end user wants to know is: "what's broken?" It is not the modem; it is not the DSLAM; it is the connection to the ISP. Debugging is the hardest and most time-consuming part of installation; this approach helps the end user quickly discover what went wrong.

Universal Plug and Play

The previous section discusses what is available to end users today. What is coming next is universal plug and play (UpnP), which should be a requirement for all future CPE from vendors because it can solve many support issues that service providers face today.

Network Address Translation
Network address translation (NAT) allows several PCs on a network to share one IP address to simplify the provisioning of services and to provide some rudimentary security services. However, NAT breaks the very applications that attract people to broadband, such as video conferencing, VoIP, multiplayer gaming, and other interactive or streaming applications. The first call that end users usually make after their broadband connection goes up is to their service providers complaining that their applications aren't working anymore.

A UPnP–enabled CPE can fix these problems. For example, with UPnP enabled on the Internet gateway, Microsoft Messenger can now configure the CPE NAT so that it listens for its calls and makes sure they are sent. An incoming call automatically gets routed to the Messenger PC. This activity is transparent to users, meaning that they do not need to configure their firewalls, open their ports, or perform other similar tasks. Also, when the application finishes, the open port automatically closes, thus reducing security risks.

Service Providers and Customer Support

UPnP will also help service providers when they want to provide new services. Customers can access, or *self-provision,* these services automatically, and they will work. Customers do not need to do anything to their PCs or networking devices to make them work with your new service. Deploying VoD or VoIP or multiplayer gaming will be much less challenging and more profitable.

UPnP also helps local-area network (LAN)–side installation by eliminating the need for software setup for applications. UPnP–enabled PCs can automatically discover and configure services deployed to the CPE without having to install software from CDs that cost money to inventory and ship and will be lost immediately by end user.

A simple example is that if the service provider changes the end user's email server address, a UPnP–enabled PC will automatically reconfigure the email application without user intervention, thus eliminating the support call that will undoubtedly follow.

Event Notification
UPnP provides event notification so that when things happen, good or bad, users can be informed. For example, suppose a wide-area network (WAN) connection goes down. Users usually call tech support and say, "I cannot connect."

Now this information can be sent to client PCs that have a CPE status display. Users will know that a WAN connection is down and being taken care of; they can also be informed of new services and asked to please call to learn more.

Service providers must imagine and plan how to implement event notification in a way that assures users that they are getting better service and have more control. Customers see service providers that notify them of problems or new services as more proactive and are less prone to switch to another carrier. When the link does go down, end users can quickly determine what the problem is. When a WAN is down, users often work through a 10-step process. They spend time and effort checking their power supplies, their cable connections, and so on. Now a user can look at the status screen. If it says the WAN connection is down, the user knows exactly where the problem is.

Event notification also makes other services in the home network easier to manage. Imagine that someone wants to disconnect their Internet connection between midnight and 7:00 am to prevent children from using it. What if another adult wanted to use the Internet during those hours? Knowing that Dad disconnected it and the LAN is not down would help that person. Notifying all PCs in the house of this event is easy, and it cures many of the headaches of managing a network.

Summary

CPE vendors can and must help ramp DSL deployment. The CPE auto-configuration efforts of the DSL Forum are a good start for WAN setup, but LAN and PC setup solutions also merit considerable attention. LAN and PC auto configuration is possible, and a standards-based solution such as UPnP does exist.

Notes

1. VPN clients are also potential customers who can benefit from UPnP, depending on their VPN software. Some VPN software that works through NAT today would not change for a UPnP Internet gateway.

Overlay Networks versus Converged Technology Platforms

Ron Longo
Vice President, Sales and Marketing
Mahi Networks, Inc.

This paper offers an in-depth comparison of overlay networks and integrated technologies, or converged technology platforms. It will draw from a number of business case studies undertaken by major carriers on the benefits of deploying converged platforms. In addition to the numerous competitive advantages that converged platforms provide, the studies demonstrate the profound capital expenditure (CAPEX) and operational expenditure (OPEX) savings that result when converged platforms are deployed in carrier networks.

Today's carriers are faced with intense scrutiny from Wall Street, along with ever-shrinking profit margins. More than ever they need technologies and business strategies that will save them dollars, bring in more income, and enhance profitability. Converged networking platforms that collapse data-communications capabilities with transport functionality meet these needs by reducing CAPEX and OPEX by orders of magnitude, supporting value-added service offerings, and enabling new, Internet protocol (IP)–based revenue models. Furthermore, converged platforms provide the added benefit of a smooth, self-paced migration path to true next-generation networking.

Historically, transport has been associated with capacity and reliability, rather than speedy provisioning. It has meant big, fat pipes, long distances, and dependable service, but certainly not dynamic switching. In fact, in the transport world, provisioning an optical carrier (OC)–3 between San Francisco and New York City can take months to complete. In contrast, data networking is all about rapid, intelligent, dynamic switching, rather than reliability. In the data world, there is not a 100 percent guarantee that a signal will get through to its destination correctly the first time. However, a person can log onto the Internet and send a packet of data from San Francisco to Budapest in a matter of seconds.

Wanting the best of the transport world and the data world, today's largest carriers are demanding converged technology. According to J.P. Morgan's Backbone! Report, the next-generation network will be a mixture of traditional transport, emerging data networks, and converged platforms—a mixture keenly designed to enable next-generation service delivery and new revenue models. The report takes care to note that during their migration to next-generation networking, carriers must not lose sight of what is "keeping the lights on"—the leased line and public switched telephone network (PSTN)–based services. Because they comprise as much as 70% of a carrier's revenue stream, these legacy services must be protected and optimized.

Securing Revenues, Growing Profits

A big challenge facing today's carriers is securing revenue while growing profits. To accomplish this, carriers must first flatten their networks. The last five years saw a huge build-out of discrete element, overlay networks. In a mad rush to win and secure customers from the growing competition that was made possible by deregulation and easily available capital, carriers of all types rapidly and perhaps unwisely "threw" equipment at their networks, building up inefficient and costly overlay structures. With the current economic downturn, these overlay networks are proving to be a dangerous drain on resources and dollars, as well as a hindrance to rapid service delivery and the deployment of new revenue-generating services. Carriers are taking a fresh look at their network architectures and are asking, "How can I optimize my existing network?" and "How can I extend its useful life while preparing my network for the coming migration to converged next-generation equipment?"

The second requirement for securing revenue and growing profits is the facilitating of new revenue models. In February 2000, carriers could charge $33,000 a month for a digital signal (DS)–3 circuit from New York to London. Twelve months later, in February 2001, the going rate was down to $12,000 a month, according to RateXchange Corporation. From a revenue-per-bit perspective, this represents a quick, tremendous decline. Clearly, carriers need a way to change this unhealthy revenue plan.

Converged technology meets carriers' needs in each of these areas. It flattens their networks while allowing them to fully leverage their installed infrastructure and extend its lifecycle. It also enables carriers to deliver advanced multimedia services and enact "revenue-per-service" or "revenue-per-value" business models.

Converged platforms also give carriers a smooth, self-paced migration path from their current infrastructures to the fully converged network architecture of tomorrow.

The Carriers' Checklist for the Next-Generation Platform

A number of common themes can be found when we analyze the requests for information (RFI) and requests for proposals (RFP) that carriers have been issuing during the past 18 months. Each of these themes point to carriers' need to drive discrete elements out of their network architectures (see *Figure 1*).

FIGURE 1

The Carriers' Checklist

RFI / RFP Requirements	Reason
Multi-Function	Flatten the network, remove elements
Multi-Protocol	Internetwork between today's and tomorrow's protocols
Terabit Scalability	Meet rising demands in the metro core
Rapid Service Activation	Faster service deployment, lower ops cost, less revenue lag
Global Platform	Meet ANSI and ETSI specs for global internetworking
Legacy Support	Full support of legacy, representing revenue majority
Non-disruptive Migration	Bridge legacy networks & Next-Gen networks, no disruption
Carrier-Class Data	Deployed in carrier networks, requires 99.999% resiliency

An examination of the RFIs and RFPs shows that carriers want multiprotocol systems capable of bringing their networks from the time division multiplexing (TDM)–based infrastructure that proliferated during the last 100 years to the IP–centric infrastructure that will be vital in the future. They want multi-terabit scalability to handle ever-increasing traffic. They want rapid service activation to secure and retain their customer base and bring revenue in-house as quickly as possible. Gone are the days when carriers can leave revenue on the table for six months or more while an OC–3 is provisioned from San Francisco to New York City. Because many of today's carriers are global by nature, they want their equipment to meet both American National Standards Institute (ANSI) and European Telecommunications Standards Institute (ETSI) specs. They need equipment that supports their current, revenue-generating legacy services. They want to provide carrier-class data services. And, they want equipment that allows them to smoothly and non-disruptively migrate from their existing infrastructure to a fully converged, next-generation architecture.

Current Architecture

Why do carriers need converged platforms? A look at the current architecture of a carrier's network provides numerous insights. Let's look at "a day in the life of an OC–48" as it travels from San Francisco to New York City (see *Figure 2*).

In a single central office (CO), the OC–48 will touch five distinct network elements (NEs), including synchronous optical network (SONET) add/drop multiplexers (ADMs), broadband digital cross-connect systems (DCSs), IP routers, asynchronous transfer mode (ATM) switches, and dense wavelength division multiplexing (DWDM) equipment. In between, the OC–48 will go through fiber-optic patch panels that are not manageable remotely. In this CO example, there are also 10 expensive, non-revenue-generating optical line interface units.

FIGURE 2

The Current Architecture

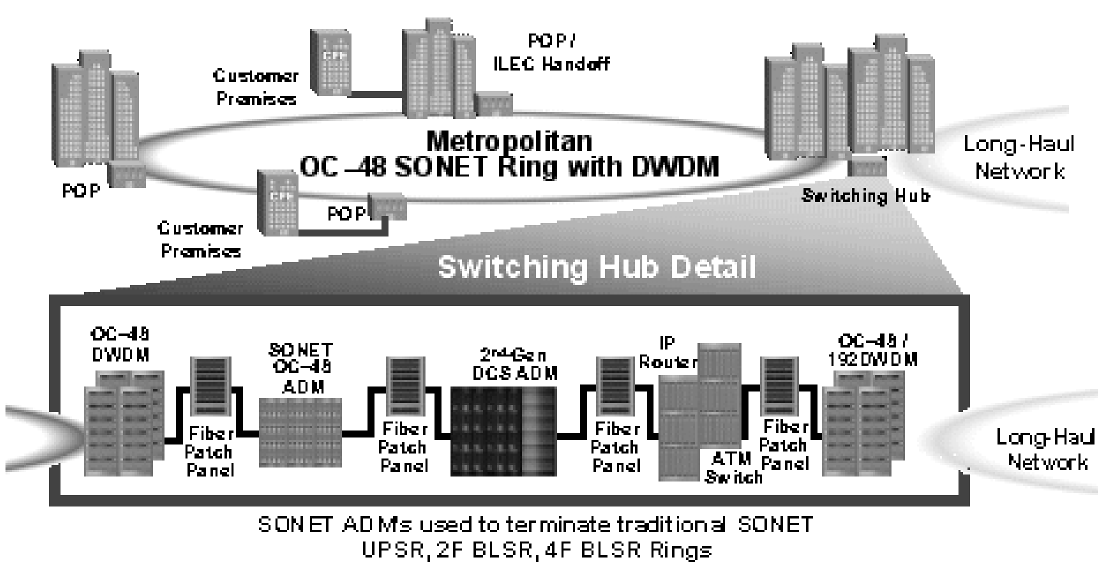

To add insult to injury, there are discreet management centers for each of these separate NEs, resulting in an incredibly complex and inherently inefficient network-management architecture to control the OC–48.

Converged Technology

With converged technology, the discreet NEs are reduced into a single platform, all managed from a unified management system. To be truly effective, converged platforms must enable "virtual separation with physical integration," which means that network operators can easily control their portion of the multiservice system via familiar interfaces. The IP manager interacts with the converged platform in the same way he or she interacts with an IP router. The broadband DCS manager interacts with the converged platform in the same way he or she has been interacting with his legacy DCS equipment. This eases the transition and lessens retraining impacts. Then, when the carrier is ready and when its business plan allows, it can smoothly migrate to a unified network operations center (NOC) for additional savings and efficiencies.

Economic Benefits

Of course, the key driver for moving to a converged platform is cost reduction—in terms of both capital and operational savings. On the far left of *Figure 5* is a depiction of a legacy network comprised of four discreet, first-generation components: a wavelength division multiplexer (WDM), an

FIGURE 3

The Current NOC Architecture

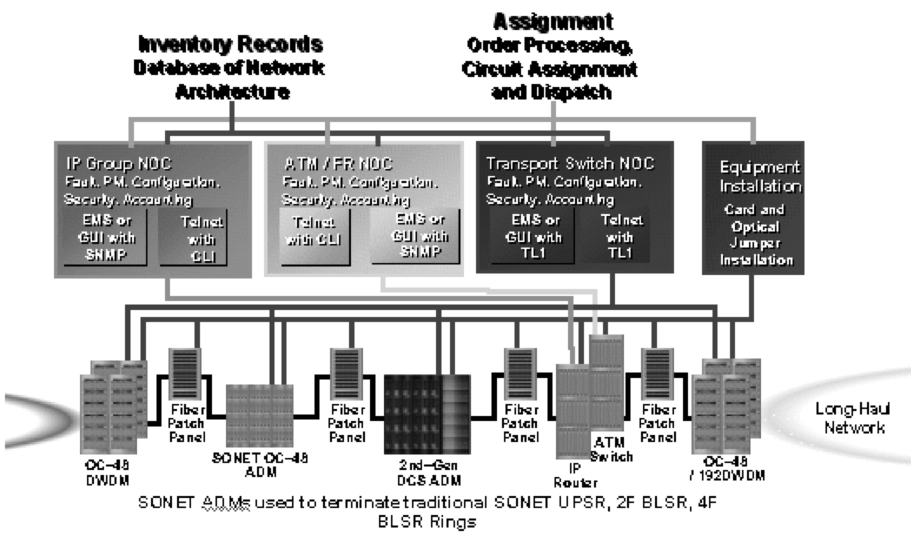

FIGURE 4

Converged Technology Platforms

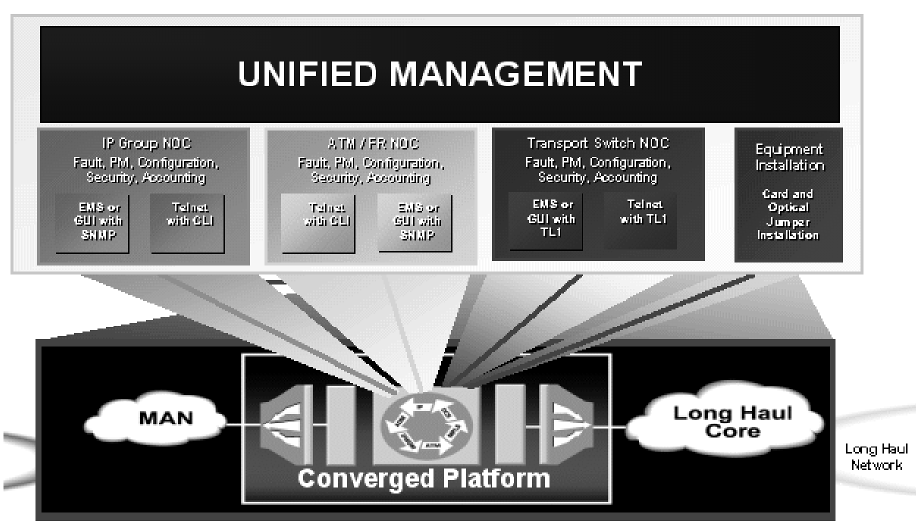

ADM, a DCS, and a router. An OC–48 going through a CO structured in this manner costs about $370,000 in CAPEX, with an additional $112,000 in OPEX.

The next column in *Figure 5* depicts the cost of the same OC–48 going through a network comprised of second-generation (2G), discrete NEs. These products take advantage of advances in application-specific integrated circuit (ASIC) technology to deliver incremental cost savings to carriers from a CAPEX perspective. However, the 2G DCS, for example, still looks like a DCS, behaves like a DCS, and is managed like a DCS. So, the network's operating expenses are not reduced, and the improvement in network functionality that they provide is only incremental.

As *Figure 5* shows, converged platforms are the breakthrough technology that deliver true, dramatic CAPEX and OPEX savings. When the network is configured through a multifunction, multiprotocol device, CAPEX savings are realized from the simple reduction in the number of elements, the reduction of physical ports from 10 to two or four, and the reduction in the number of non-revenue generating optical line interface units. OPEX savings are realized as space and power requirements are reduced, and as truck rolls are cut back. In all, numerous case studies conducted by some of the world's largest carriers have shown an average savings of more than 60% in CAPEX and OPEX through converged technology platforms.

FIGURE 5

The Economic Benefits

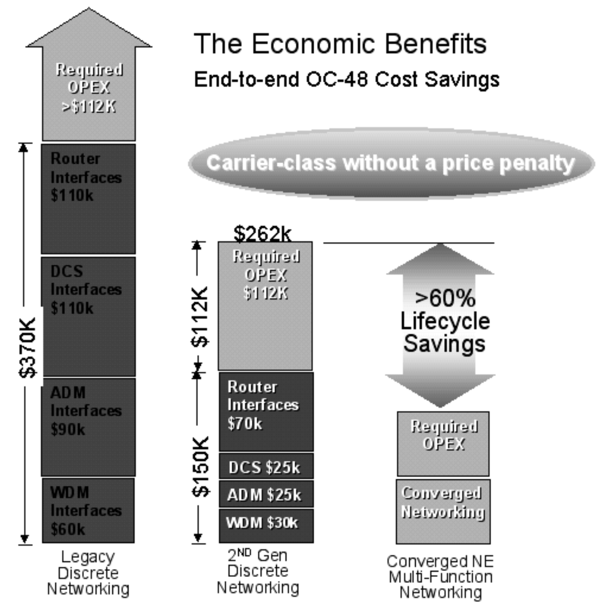

Conclusion

In summary, market conditions are forcing today's carriers to focus on their bottom lines and to implement more intelligent business strategies. Quite simply, today's discrete element, overlay networks are incapable of meeting the growing demands for bandwidth and next-generation services, as well as the scrutiny of today's financial markets. Converged technology platforms give carriers a "smarter" network—one that supports and maximizes existing revenue streams with zero disruption to current operations, reduces capital and operational expenditures by orders of magnitude, and provides an easy, "at-their-own-pace" migration path to a next-generation network topology capable of rapidly delivering higher margin value-added multiprotocol services.

Next-Generation Optical Metro Networks and the Global Lambda Grid

Joel Mambretti
Director, International Center for Advanced Internet Research
Northwestern University
Director
Metropolitan Research and Education Network

Introduction

Even as the digital communications industry develops new economic models in response to a more dynamic market, new opportunities are being generated by next-generation networking technology innovation, which continues to advance at an ever-faster pace. Not only are multiple next-generation Internet research and development (R&D) projects and implementations progressing quickly (MAM99), but leading-edge infrastructure architectures, technologies, and methods are also enabling powerful new applications and services, particularly those based on optical networking (STE99, RAM99). These developments are particularly observable in metro areas, as optical networking technologies traditionally used in long-haul networks migrate to regions, cities, and enterprises. Optical networks are becoming a foundation for digital communications services not possible previously, such as those based on dynamically provisioned wavelengths. This technology transition is giving rise to a more flexible transport model for networks, providing for services based on extremely high performance routers interlinked with intelligent core optical networks. This approach can support multiple new services with digital transparency. Some of these new designs are now beginning to be realized in prototype infrastructure and services, including Internet protocol (IP)–over-lightwave networks, controlled by IP signaling. Such signaling can be used not only for the dynamic provisioning of lambdas "on demand," but also for other types of rapid resource allocations. Within metro areas deploying technologies based on lambda switching, it will soon be possible to create OCX facilities as exchange points for "global lambda Grids," utilizing optical VPNs. These global Grids will become a powerful platform for multiple, innovative distributed worldwide applications.

Architectural Design

Emerging digital services and high-performance applications require significantly enhanced network architecture. Infrastructure designed to support optimally traditional services, such as circuit-switched voice communications, is being replaced by infrastructure that is optimized to support connectionless data communications—digital flows that are less predictable, characterized by bursty traffic, asymmetric streams, and multiple temporal and local variations. For example, increasing high volumes of self-similar Internet traffic, as well as a wide range of bandwidth-aggressive next-generation applications (some individually requiring 10 gigabits per second [Gbps] point-to-point bidirectionally), cannot be supported appropriately by current architectural implementations. Consequently, standards bodies are developing architecture for new, flexible, scalable, and high-performance digital communications, and these designs are beginning to be realized in prototype infrastructure and services. For example, the Internet Engineering Task Force (IETF), the Internet architecture standards organization, has a concept of a powerful, "smart" network core, while placing at the edge of the network as much complexity and intelligence as possible. In contrast, the classic telecommunication model was oriented to provisioning complexity at the core, e.g., within central offices (COs) and simple devices at the network edge.

Many current architectural initiatives are directed at removing hierarchical layers to eliminate cost and complexity and to advance the goal of providing digital services with transparency. Although today some data services are still supported by an infrastructure consisting of asynchronous transfer mode (ATM) over synchronous optical network (SONET)/synchronous digital hierarchy (SDH) over fiber, other providers have eliminated the ATM layer and in favor of packet over SONET (POS). The next step is to remove the SONET layer. Although SONET is highly reliable and provides for efficient, well-known management and traffic-engineering techniques, it is neither optimal nor cost-effective as a support foundation for digital services. Its equipment and management costs remain high and its typical deployment in two parallel bidirectional rings—one for production, the other for protection—requires 50 percent of deployed capacity to be idle. Also, it does not allow for fast deployment of new and enhanced services and capacity. To replace the reliability and restoration capabilities of SONET, a range of techniques can be used in conjunction, including implementing multiple potential paths via mesh topology and using standard and extended L3 IP techniques. Some form of framing is still required, and several potential options have been suggested for this function, such as lightweight SONET, Ethernet, and new types of protocols. At this time, Gigabit Ethernet (GbE) is a particularly popular framing techniques, driven in part by the growing demand for metro GbE services as enterprises extend local-area networks (LANs) across cities.

To a significant degree, this architecture takes advantage of the capabilities of the same dense wavelength division multiplexing (DWDM) and related technologies that have been used in long-haul networks for since the late 1980s and that have begun to be deployed in regions and metro areas. DWDM divides a beam of light into multiple "colors," lightwaves, or "lambdas," so that multiple optical signals can be transmitted through fiber in parallel—extremely high volumes of data can be communicated through each strand of optical fiber. Newer technology provides capabilities for dividing the light into many dozens of channels on each fiber length and for sending tens of gigabits or more through each channel. Single fibers are capable of transmitting extremely high volumes of data, which will lead to ubiquitous broadband and significantly lower cost for digital services. The many lightwave channels within DWDM–based networks not only provide increased capacity for high volumes of data traffic within a common physical infrastructure, but they also allow different wavelengths to be dedicated to a separate services and classes of service. Multiple types of digital communications services based on different classes can be provisioned. Designing and redesigning networks by manipulating lightpaths instead of creating and dismantling physical infrastructure allows for flexibility when creating digital communications systems. Optical networks also have a significant potential for rapid direct provisioning and redeployment of individual wavelengths and services on those wavelengths, a feature sometimes called "point-and-click provisioning."

Because these techniques can be implemented more easily in metro areas, some of the earliest deployments may be within cities, supported by a new generation of optical metro devices. Many first-generation optical metro devices were merely long-haul systems that were redeployed within cities, for example, to support transit. Newer devices are being designed specifically to meet the requirements of metro-area digital communications services. In addition, innovative tunable filters and tunable lasers are being incorporated to allow enhanced flexibility and customization. Other efforts are being directed at creating complementary metal oxide semiconductor (CMOS)–based OCXs as edge and enterprise devices. Consequently, in a number of metro areas, consideration is being given to transitioning from SONET over time and to provisioning digital services directly over fiber, creating a new type of optical network—one that is less oriented to point-to-point implementations and more to the dynamic provisioning of digital services based on lightwaves ("lambdas on demand").

Signaling Architecture

There are various approaches to implementing signaling architectures related to provisioning digital services over optical networks. Some current efforts are directed at defining an IP control plane that would interface with optical control and traffic planes. There are four general models, with multiple hybrids. One is an overlay model, which defines separate domains. Another is IP–centric, and yet another is optical-centric. In this model, IP transit is supported by optical networks, much as it was supported by ATM. In this model, lightpaths are much like permanent virtual circuits (PVCs), and fairly static provisioning is presumed.

Another is a signaled overlay model, which provides for separate domains, but also for a wide range of interactions between the IP layer and the optical layer, allowing for resources on demand, end-to-end lightpaths between any two edge points, via out-of-band signaling. This architecture, which is the focus of significant development, provides for signaling for various levels and types of services and related requirements (e.g., priority, protection class, etc.), resource discovery and availability mechanisms, lightpath management (create, delete, change, swap, reserve), optimization and other performance processes, survival/protection and restoration processes, etc. Another architectural model is the peer model, which envisions IP devices (e.g., routers) and optical devices (optical switches) as cooperative components, each with IP identifiers and within a shared infrastructure fabric and using a common routing protocol and state data. A fourth major model is an integrated model, which envisions a single network fabric with a high degree of integration of IP and optical components.

Architecture Standards Activities

The International Telecommunication Union–Telecommunication Standardization Sector (ITU–T) is defining architectures for an optical transport network and automatic switched optical networks, including provisions for network-to-network interface (NNI), user network interface (UNI), control signaling, etc. The Optical Internetworking Forum (OIF) is also developing architectures for UNI and NNI. The Optical Domain Service Interconnect (ODSI) forum is developing architectures defining UNI specs for routers as clients to optical systems, and OXCs (all optical as well as electric core OCXs). They are also beginning to address avelength division multiplexing (WDM) line terminals and optical add/drop multiplexers (OADMs). The IETF has undertaken a variety of efforts within multiple working groups. Key efforts are those that relate to the link management protocol and the control plane efforts related to GMPLS.

IP–Centric Control Plane Architecture for Optical Networks

As the progress toward IP over optical networks continues, the importance of control-plane architecture has been increasingly recognized. The IETF has undertaken a number of efforts related to designing a common IP–centric control plane that would provide for precise out-of-band management of lightpaths within and across domains. The development of these management capabilities will allow for greater network flexibility, for example, replacing static point-to-point WDM implementations with methods for dynamic provisioning of not only lightpaths, but also a wide range of other network resources. Consequently, investigating optimal architectures for such a control plane is an important issue—especially those providing extensions related to new technologies, such as enhanced photonic components, all-optical cross-connects, channel distance attributes, and large numbers of high-performance channels. An IP control plane must provide for extremely high performance capabilities for a variety of functions, such as optical node identification; service-level descriptions (e.g., request characterizations); managing link-state data, especially for rapid revisions; allocating and re-allocating resources; establishing and revising optimal lightpath routes; determining responses to fault conditions; etc. General functions include the following:

- Specifying and updating lightpath addressing, employing unique identification of path end points
- Determining lightpath availability and reachability
- Dynamically provisioning lightpaths through lambda processing (discovery, add, delete, switch, change, restore, etc.)
- Multiservice and multiprotocol supports, including IP, GbE, 10 GbE, etc.
- Traffic engineering
- Performance monitoring and analysis
- Survival and protection

GMPLS

Currently, the generalized multiprotocol label switching (GMPLS) architecture (MAN01, BEL01, ABO01, ASH01a), being developed by the IETF, has gained significant momentum. The primary IP–over-wavelength optical architectural models previously described and can utilize GMPLS. GMPLS is the generalized extension of multiprotocol label switching (MPLS) (CAL99, ROS99), a signaling protocol with a flexible framework which specifies separating forwarding information from IP header information, allowing for forwarding through label swapping and various routing options. The MPLS architecture assumed a forwarded plane that recognized packet or cell boundaries and provided processes based on packet or cell headers. GMPLS provides for extensions that include forwarding planes that are not capable of recognizing such boundaries, such as traditional devices based on time division multiplexing (TDM) (e.g., SONET add/drop multiplexers [ADMs]) and newer devices, based on wavelengths (optical lambdas) and spatial switches (e.g., in flow port or fiber to out flow port or fiber). Consequently, GMPLS allows for forwarding decisions to be based on time slots, wavelengths, or ports. Path determination and optimization is based on labeled-switched path (LSP) creation and implementation. This process gathers the information required to establish a lightpath and determines characteristics, including descriptive information (address identifiers, reachability, etc.)

Extensions to GMPLS provide for a range of extensions for routing, open shortest path first (OSPF), and general traffic-engineering functions. GMPLS–TE extensions include those that allow for constraint-base routing–label distribution protocol (CR–LDP) specific formats and mechanisms and for resource reservation protocol–traffic engineering (RSVP–TE) signaling (ASH01b). Path protection ensuring the protection of existing paths is a key requirement, and requires continual high-performance monitoring of state information. Detecting and locating faults at both the IP and optical layers and rapid responses are also high-priority functions. Currently, GMPLS development is focused on intra-domain networks. Other efforts with the IETF are attempting to develop techniques that address interdomain networking while avoiding link-state protocols and their related complexity (STA00, XUY01).

Given the importance of the IP control plane, the high level of activity, the multiple emerging standards, the amount of current architectural tasks, and the controversies, the IETF has formed a specific effort, the Common Control and Measurement Plane (CCAMP) working group, which is coordinating various efforts defining a common control plane and a separate common measurement plane.

Physical-Layer Considerations

An important consideration related to advanced optical networks based on lambda switching is that lower-level functions require ongoing monitoring of various physical processes to protect multiple potential impairments, including conditions resulting from impairment accumulation. Physical performance monitoring, verification, adjustment, and restoration are required for such functions as output leveling across wavelengths, excessive decibel (dB) loss, maintenance of wavelength precision, chromatic dispersion effects, etc.

Advanced Research Projects

Networks based on wavelength provisioning are rapidly migrating from long-haul service domains to regional areas, metro areas, and the enterprise. To take advantage of this trend, the International Center for Advanced Internet Research (iCAIR), with partners worldwide, has established a number of major R&D projects that are developing new capabilities based on wavelength switching. These projects include those that are creating high-performance applications for network-linked computer clusters (using optical links as computer backplanes), new types of network middleware (IETF RFC 2768, AIK99), and next-generation optical network infrastructure. With the Electronic Visualization Lab of the University of Illinois, iCAIR is designing a global optical exchange point (StarLight). iCAIR is also participating in the architectural design and development of international, national, regional, statewide, and metro-area optical networks, including one that will connect as many as 2,000 sites in Chicago. Recently, iCAIR established its Optical Metro Network Initiative (OMNI), which is creating a reference model for next-generation metro networking infrastructure.

References

1. (MAM99) Mambretti, J. and Schmidt A., "Next Generation Internet: Creating Advanced Networks and Services" (Wiley Press, 1999).
2. (STE99) Stern, T. and Bala K., *Multiwavelength Optical Networks: A Layered Approach* (Addison Wesley, 1999)
3. (RAM99) Ramaswami, R. and Sivarajan, K., *Optical Networks: A Practical Perspective* (Morgan Kaufmann, 1998)
4. (MAN01) Mannie, E., et al., *Generalized Multi-Protocol Label Switching (GMPLS) Architecture* (2001), draft-ietf-ccamp-gmpls-architecture-00.txt.
5. (BEL01) Bellato, A. et al., G.*709 Optical Transport Networks GMPLS Control Framework* (2001); Aboul-Magd, O., et al., A *Framework for Generalized Multi-Protocol Label Switching (GMPLS*) (2001) draft-many-ccamp-gmpls-framework-00.txt.
6. (ASH01a) Ashwood-Smith, P. et al., *Generalized MPLS – Signaling Functional Description* (2001), draft-ietf-mpls-generalized-signaling-06.txt.

7. (CAL99) Callon, R, et al., "A Framework for Multiprotocol Label Switching" (1999), draft-IETF-mpls-framework-03.txt.
   (ROS99) Rosen, Eric, et al., "Multiprotocol Label Switching Architecture" (1999), draft-ietf-mpls-arch-05.txt
8. (ASH01b) Ashwood-Smith, P. et al., *Generalized MPLS Signaling – CR–LDP Extensions* (2001), draft-ietf-mpls-generalized-cr-ldp-04.txt; *Generalized MPLS Signaling – RSVP–TE Extensions* (2001), draft-ietf-mpls-generalized-rsvp-te-05.txt.
9. (STA00) St. Arnaud, B., Weir, M., and Coulter, J., "BGP Optical Switches and Lightpath Route Arbiter," Whitepaper (2000).
10. (XUY01) Xu, Y., *A BGP/GMPLS Solution for Inter-Domain Optical Networking* (2001), draft-xu-bgp-gmpls-01.txt.
    (AIK99) Aiken, Robert et al., IETF Middleware RFC 2768

Enabling In-Service Migration from Ring to Mesh Architectures in Metro Networks

Emanuel Nachum
Vice President, Marketing
Lightscape Networks

Introduction

The demand for bandwidth for multiple services has been increasing at lightening pace. As new network applications such as e-commerce, network storage, and corporate data and video continue to grow, service providers are struggling to grow the bandwidth in step with the traffic. While optical technologies can provide the bandwidth for these high-traffic applications, one of the biggest challenges facing the industry today is building a network that meets today's needs but scales in-service to meet those of future years.

To accomplish this, network flexibility in terms of supporting multiple network topologies, such as ring and mesh, is a critical factor. Choosing multi-topology, multilayer integrated platforms, which can migrate easily from ring to mesh architecture in-service, will assure the flexibility and scalability needed to allow carriers to scale networks cost-effectively on a "when and where" basis.

Until now, metro optical transport networks, built during the last 10 years, were based primarily on synchronous optical network (SONET)/synchronous digital hierarchy (SDH) equipment deployed in a hierarchy of ring topologies as access rings, collector rings, and metro core rings. Such metro networks were appropriate for the voice era, when carrier traffic was highly predictable and slow to change.

Most analysts agree that data traffic will dominate over voice in our networks from now on. With the rapid growth in data traffic, these metro SONET/SDH ring hierarchies have become a scalability and provisioning bottleneck in connecting users and businesses with core, national, and regional backbones. Additionally, the ring hierarchy-based networks are rigid time division multiplexing (TDM) pipes and require up to 50 percent capacity for protection. These networks waste the majority of bandwidth for traffic engineering, upgrades, and protection mechanisms, especially for the dynamic nature of data traffic. For example, an optical carrier (OC)–12 metro ring can be upgraded to no less than an OC–48 ring, which implies significant waste if the user needed only marginal increase over the OC–12 bandwidth. An entire ring has to be upgraded to increase bandwidth even in a section. Additionally, any geographic extension of the network to include extra nodes requires an additional ring (or rings) to be built. Gradual extensions, which include one node at a time with point-to-point links, are impossible.

To alleviate these issues, many service providers are looking to transition from ring- to mesh-based metro network architectures. Mesh architectures are flexible, bandwidth-efficient, and cost-effective, and offer on-demand scalability and tiered-protection services. Additional bandwidth can be turned up on demand even on a link-by-link basis. Adding new sites requires mere point-to-point fiber extensions. Depending on the application, a choice among multiple protection services can be offered to customers. For example, customers can choose to leave their traffic unprotected or choose to protect with various levels of SLAs, including guaranteed protection switching in the sub 50 ms range for mission-critical applications. Such differentiated services are likely to be critical factors needed to maintain a competitive advantage and develop new revenue streams.

Most national carriers have been building their dense wavelength division multiplexing (DWDM)–based core backbone networks in mesh topology. The next step in this evolution is going to happen in the metro networks. In the metro networks, carriers need network solutions that support today's ring topology while providing migration to mesh—building it section-by-section if customer traffic demands warrants. To meet these needs, carriers should be able to configure (or reconfigure) in-service to multiple topologies in the network, including ring, star, linear, and mesh or virtual mesh-in-ring.

Carriers can build their networks today with next-generation hybrid optical platforms that support both multi-ring SONET/SDH architecture as well as DWDM–based optical mesh architecture. Carriers can initially run these metro networks as SONET/SDH rings and later migrate in-service to optical mesh, allowing them to scale their networks and/or add differentiated services.

This paper covers the following:

- The benefits and challenges of running current metro networks with SONET/SDH ring architectures
- Suggestions on how next-generation platforms can support migration from today's SONET/SDH ring architecture to DWDM–based optical mesh architecture
- Descriptions of the advantages of running mesh-based metro networks

Challenges with Rings Architectures

Current metro optical transport networks are based primarily on SONET/SDH equipment, deployed in a hierarchy of OC–3/OC–12/OC–48 ring topologies. Among the key benefits of SONET/SDH hierarchy are self-healing and survivability, features that have served the carriers well for voice traffic and cannot be compromised in other applications. SONET/SDH ring networks are relatively immune to traffic interruption because of the dedicated multiple protection paths provided in the ring itself. SONET/SDH rings can recognize a fiber cut and reroute traffic along the alternative pre-allotted protection path, before any significant degradation in performance can occur.

There are two types of SONET/SDH rings: a unidirectional path-switched ring (UPSR) and a bidirectional line-switched ring (BLSR). The former (Bellcore spec TR-496, path switched) switches individual paths and the latter (Bellcore spec GR-1230, line switched) switches the entire optical line capacity. There are differences in their architectures that significantly affect the bandwidth utilization of the network.

Major carriers in the United States have typically built their metro networks as SONET/SDH rings. UPSR rings are commonly found in the metro access and aggregation rings, while BLSR rings are more often deployed in the metro regional core network and sometimes even the national core network. Additionally, most carriers have typically used linear or mesh connectivity for their long-haul inter-city DWDM networks. Even the SONET/SDH traffic on these long-haul networks may be in point-to-point configuration, which is unable to capitalize on the survivability features of the SONET/SDH rings. As a result, cable cuts in these sections could result in an outage for voice traffic, unless protected by an alternative protection mechanism in the optical layer. However, data-networking devices, such as Internet protocol (IP) routers, do not necessarily need SONET/SDH or optical protection, as they can usually identify the broken link and reroute the traffic around the fiber cut sections.

UPSR sends traffic both ways around the two-fiber ring for redundancy, thereby reserving 50 percent bandwidth for protection. The receiving end monitors both signals and selects the better one. Because of this dedicated use of protection capacity, UPSR systems have less fundamental capacity for paying traffic than BLSR rings. A four-fiber BLSR ring sends traffic only in the required direction during normal operation. Each direction has a working fiber and a protected fiber. During a fiber interruption, the traffic is routed around the break in the opposite direction. This method allows the ring to share the protection capacity, which increases the overall traffic carrying capacity of the system.

Ring architecture has proven very reliable for sub 50 ms protection guarantees that were important for the networks carrying predominantly voice traffic. BLSRs are very effective when the metro core is a single SONET/SDH ring, interconnecting multiple central offices (COs) or points of presence (POPs). Indeed, BLSR performs three functions very well. First, it allows the network to recover from failures very quickly, normally in less than 50 ms. Second, it provides a shared pool of protection bandwidth within a given ring so that carriers do not need to pre-allocate protection bandwidth on a one-for-one basis. Third, it allows pre-emptable traffic to traverse the protection channels when they are not carrying protection traffic.

However, the BLSR begins to show limitations as traffic growth forces service providers to install multiple parallel BLSRs. Because a BLSR only provides the ability to share protection traffic within a given BLSR, multiple parallel BLSRs cannot have a shared pool of protection bandwidth amongst themselves. Also, a BLSR is required to be symmetric, with the same amount of working and protection bandwidth all the way around the ring. These limitations result in the addition of "extra" or "wasted" bandwidth whenever the demand on the various segments of the SONET/SDH rings is unequal. With the BLSR, each ring must be sized to accommodate the busiest segment of that ring. The inefficiency of this solution is apparent by looking at the amount of working bandwidth on each segment of a ring in comparison to the traffic on other segments. Furthermore, the inability to share protection between adjacent rings adds unnecessary segments. Clearly, the efficient restoration technique of the BLSR ceases to be bandwidth efficient while scaling the metro core network to a multi-ring environment.

To scale the bandwidth of any section in a ring, the entire ring has to be matched up with that capacity. This leads to significant waste of network capacity. The problem with UPSR/BLSR architecture goes beyond the wasted protection bandwidth. Carriers lose competitive advantage when they are not able to carry out on-demand network scaling cost-effectively. Moreover, the BLSR does not enable service providers to define differentiated protection services.

In summary, the BLSR, when used for multi-ring metro core or regional/national backbone network restoration, is not the most efficient way of using the network bandwidth. Additionally, it limits the carriers' ability to offer differentiated services and cost-effectively scale their network in-service. In contrast, an optical meshed architecture eliminates these deficiencies and makes network provisioning, operation, and scaling highly cost-effective, providing next-generation carriers with a significant competitive edge. In the absence of support for mesh architecture in their backbone network, carriers sometimes create a set of interconnected rings (logical rings over physical mesh) to cover regions and even an entire country. This helps carriers provide at least for the survivability of the network without adding any of the advantages of the mesh network previously outlined.

Carriers should begin building networks that support SONET/SDH rings *and* allow them to migrate cost-effectively to optical mesh architecture in-service at a later date.

Migration to Mesh Architecture

Mesh optical networks are the future model for all networks, both metro and national backbone. Carriers need to build ring-based networks today, especially in metro areas, and need to be able to migrate in-service to mesh in the near future. Carrier's next-generation networks can support both ring and mesh topology at every node today. Topology definitions of their network are embedded in the element-management system (EMS) and network-management system (NMS) and do not require replacement of any hardware at the time of migration from ring to mesh architecture. In this process, if the traffic routes are to be modified and higher traffic to be accommodated in certain links, the required hardware changes also can be undertaken in-service. Such an upgrade is carried out in three easy steps: (1) route the live traffic over its protected path momentarily, (2) replace the card, and (3) route the live traffic back on to the upgraded card.

Metro mesh networks support not only the survivability attribute of the ring architecture in the optical layer, but also offer multiple levels of survivability options based on dedicated or shared protection bandwidth or dynamic search for protection routes (e.g., Gold, Silver, and Bronze classes). Additionally, it adds another distinct value to a metro optical network: instant and cost-effective scalability of the network.

The network shown in *Figure 1* is a representation of how next-generation multi-topology platforms can recognize both SONET/SDH rings and optical mesh architectures in a network. This allows seamless migration from ring-based networks to mesh. As a mesh, this network can grow not only the bandwidth of any of its existing links in-service, but also allows for physical extensions of the network by building a new point-to-point link (or links).

Advantages of Mesh Networks

Survivability with Bandwidth Efficiency

In mesh networks, service provisioning between two points takes the traffic along the most cost-effective path available with tiered protection, as desired and paid for by the customer. Depending on the type of service, it could mean allocating a predefined reserved 1:1 protection path (Gold service), 1:N shared protection path (Silver service), or a dynamic routing algorithm to search for a new restoration path (Bronze service). Mesh networks can provide transport service without compromising the guaranteed protection of sub 50 ms for mission-critical applications required by banks, financial institutions, and real-time e-commerce traffic. However, for many corporate data services, this architecture can allow carriers to offer no pre-allocated (and wasted) protection bandwidth. In such cases, the network recovers from a breakage by using the intelligent automatic fault sensing and rerouting feature of its optical switches. Assuming many of the carriers' services are likely unprotected or shared protection data services, the potential saving of network bandwidth is dramatic. This adds significant cost advantage to the carriers and helps them to beat the competition.

FIGURE 1

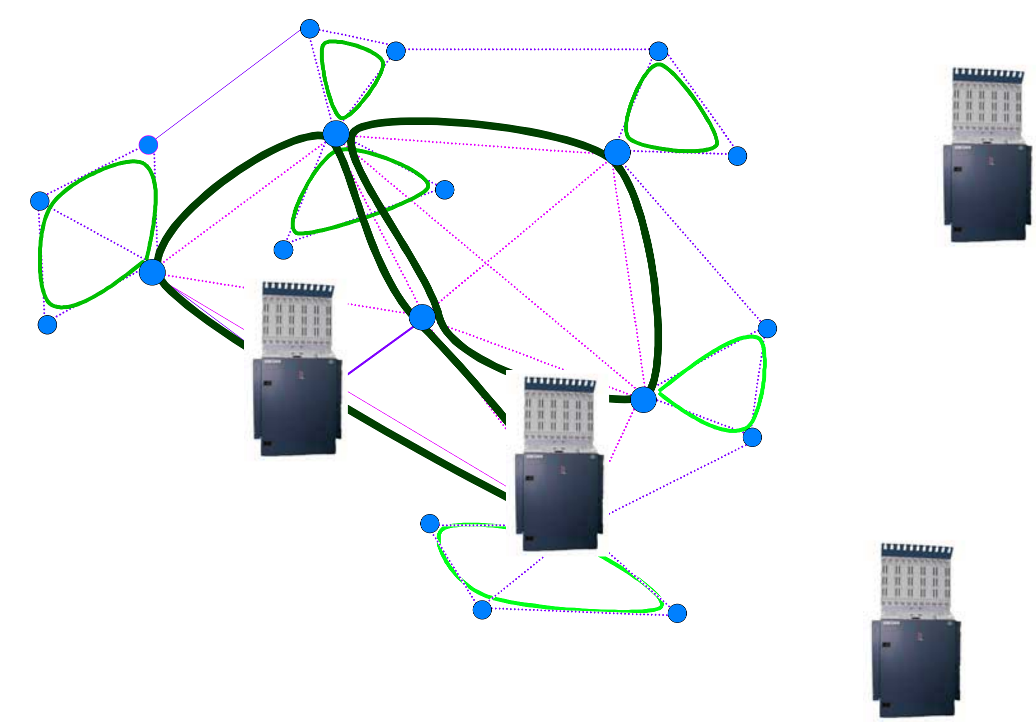

Based either on dynamic routing algorithms or pre-allocated protection bandwidth, an optical switching-enabled mesh architecture can also provide self-healing capabilities without the complexity associated with ring software. Some intelligent optical switches are designed with software intelligent enough to retain the advantages of the BLSR while at the same time eliminating the inefficiency of multiple parallel rings. Introduction of optical switching transforms the network topology from a ring or linear architecture to a mesh architecture, providing for optimum routing of lightpaths from source to destination and fast rerouting around network failures without the need to reserve protection bandwidth for all types of applications.

As an example of bandwidth savings under predefined shared protection scheme (see *Figure 2*), consider 1:4 shared paths as follows:

- Three different optical trails are routed between sites A and D (full lines).
- One reserved shared path (dotted line) offers protection bandwidth for the four working paths, allowing for a significant bandwidth savings over a 1+1 protection that wastes 50 percent of the network.
- Upon sensing a failure in one of the working paths, traffic switches to the shared path automatically by the optical layer.

Revenues from Tiered Protection Services

One of the challenges in today's climate is to help carriers maximize their revenues by evolving from just selling bandwidth as a commodity to providing more enhanced services. To face this challenge successfully, carriers need to offer both differentiated premium services as well as basic services. Mesh architectures enable a carrier to sell tiered services, such as Gold (guaranteed), Silver (guaranteed with delays), and Bronze (best effort). This strategy can help a carrier to generate new revenue streams by differentiating from others that may be unable to sell any premium services.

Data-Optimized Network

It is known that overlaying data services networks, such as IP and asynchronous transfer mode (ATM) backbones, are mesh networks, made up from complex interconnections of IP/ATM routers and switches. These devices choose the most efficient of the multiple options over the mesh to ship traffic from origin to destination. Naturally, traffic engineering of IP and ATM traffic is more aligned with metro optical mesh networks. This facilitates better connectivity and bandwidth utilization throughout the network. Being bitrate and format independent, the optical mesh can carry both data and voice traffic efficiently.

The emergence of optical cross-connects (OXCs) represents a landmark in evolution of mesh networking, as it enables current networks built up with point-to-point DWDM links and SONET/SDH rings to evolve into a unified next-generation intelligent mesh network. The DWDM/optical layer not only provides hundreds of channels of capacity for the explosive growth of data traffic, but also becomes a means to integrate multiple data services on to an intelligent mesh network. Alternate routing through the mesh either by the

FIGURE 2

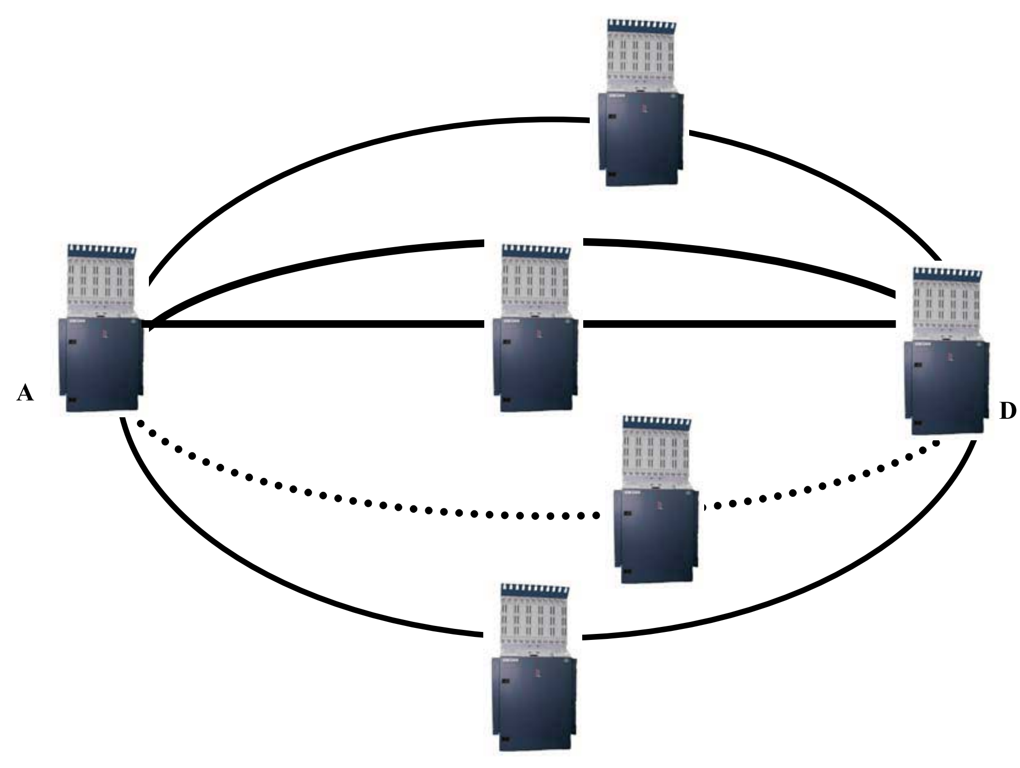

IP/ATM router/switch or OXC becomes a highly optimized solution that can match tired protection needs of multiple applications. Additionally, metro mesh networks also enable carriers to offer auto-provisioned lambda and data services efficiently.

Network Scaling

As the number of lightpaths or channels traversing the network increases, optical switching in a mesh environment is required to continue to efficiently scale the core of the metro network. Mesh networks can truly scale on a "when and where" basis: carriers can add capacity to existing networks on a link-by-link basis—*on demand*—without having to upgrade an entire ring or network. They can also extend the geography of the network to new nodes by merely adding point-to-point links that become part of the single unified mesh network, instantly and effortlessly.

Conclusion

Ring topology-based metro networks served the carriers well in the voice era, when a network-wide self-healing in the sub 50 ms range was the key issue. Limitations of rings toward bandwidth efficiencies and on-demand scaling were not critical. Now, carriers need to build metro networks that support both ring and mesh architectures and can migrate from ring to mesh cost-effectively and in-service.

Mesh-based metro networks are more bandwidth efficient, as there is not always a need to reserve additional spare capacity for protection purposes. The tremendous cost advantage of mesh networks is the ability to free up to 50 percent of bandwidth for revenue-generating traffic, providing carriers with a significant competitive edge.

Mesh-based networks enable service providers to configure tiered services. That means providing only as much working and protection bandwidth as is actually being paid for and needed to support the contracted service-level agreement (SLA). For example, mission-critical services can receive 100 percent guaranteed sub 50 ms protection, while non-critical services can receive shared, delayed, or even no protection at all. Only selected (paying) services need to be provided with additional spared protection bandwidth, not the entire network.

Mesh networking provides carriers with on-demand and in-service scaling capabilities. Carriers can add bandwidth only in the sections or links where additional traffic demands.

HomeRF: Optimized for the Broadband Internet Home

Kevin J. Negus
Chief Technical Officer
Proxim

This paper will discuss technology from the Home Radio Frequency Working Group (HomeRF). HomeRF leverages what digital subscriber line (DSL) brings to the household. This paper will introduce the broadband Internet home, discuss some of the candidate wireless technologies for networking that home, and then go into detail on how HomeRF works. It will end with a comparison of candidate technologies.

Broadband Internet Households

A lot of homes in America are going to have broadband services in the near future. *Figure 1* shows the projected numbers of households that will have broadband services such as cable, DSL, and wireless by 2004.

Services Consumers Want from a "Telecom Bundle"

What all U.S. consumers expect from a bundled set of services coming into the broadband Internet home—the one feature that stands out—is voice (see *Figure 2*). But most people at networking or data conferences do not believe it—they believe high-speed Internet is the driver of everything. And while it certainly is a big factor, voice, streaming media, audio, and video are all equally or more important in terms of bringing revenue into the service provider.

Candidate Wireless Home-Networking Technologies

The large number of wireless networking technologies can be very frustrating. There are definitely many more than the three discussed here. The bad news, depending on one's perspective, is that there are many more to come. The nature of wireless is that it is an unbounded medium. It is not a physical wire. The physics of wireless enable more consumer choice, so consumers demand and get an optimum solution for whatever application they have.

Consequently, wireless equipment providers cannot have one wireless standard that does everything. It would end up being substandard in performance for any given application and probably would not sell. Thus the industry ends up with a lot of different standards, and there is no easy alternative.

What makes wireless particularly difficult is not only that there are a lot of different standards, but that they interfere with each other. They go beyond not working together and actually hurt each other. That is definitely true of the three discussed here, which all operate in the 2.4 gigahertz (GHz) band.

IEEE 802.11b (or "WiFi")

The Institute of Electrical and Electronics Engineers (IEEE) 802.11b technology is one of the latest in a long line of wireless technologies going back to OpenAir, WaveLAN, and the original 802.11DS and 802.11FH. All of these, in different ways, are trying to be the wireless equivalent of Ethernet.

HomeRF

In terms of data, HomeRF is in many respects no different from 802.11b. It is wireless Ethernet as well. But it adds the toll-quality cordless voice capability that Digital European Cordless Telecommunication (DECT) has, which has been the most successful digital multiuser cordless standard in the world to date. HomeRF also adds the ability to do multiple prioritized streaming-media sessions, which currently are typically audio or low-rate video sessions. Its other advantage compared with 802.11b is that it addresses the major interference and security problems inherent in that technology.

Bluetooth

The third technology, Bluetooth, has had a lot of fanfare and actually has some tremendous potential. Bluetooth is also a 2.4 GHz wireless communications technology, but it provides a unique wireless connection that is literally always on and always able to make an ad hoc connection, yet draws very little power. Of course, solving that problem does not necessarily make Bluetooth a backbone home network, a wireless version of Ethernet to the home, or a cordless telephone. It was really not designed for networking, but rather as an extremely low-cost and low-power cable replacement.

Nevertheless, there are tremendous predictions for Bluetooth penetration, including it being in every cell phone

FIGURE 1

Broadband Internet Households

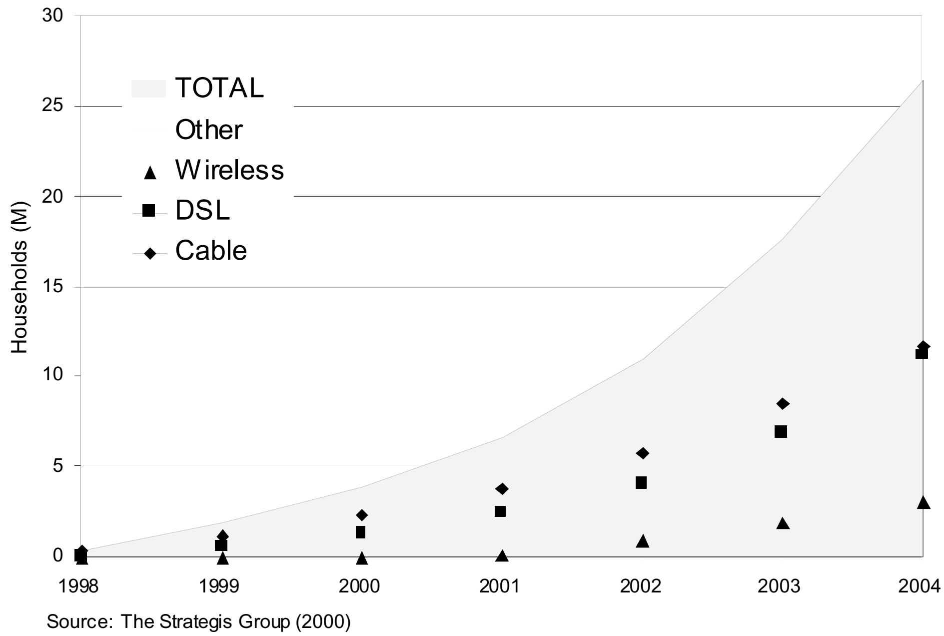

and personal digital assistant. And if Bluetooth is going to be in these devices, it will be in the home. Bluetooth is therefore a major factor for home wireless networking, even if just as an interference source. Being able to operate in the presence of multiple Bluetooth devices is important for any successful home wireless-networking technology.

HomeRF 2.0 Capabilities Summary

Version 2.0 of HomeRF, which is shipping today, has a 10 megabit-per-second (Mbps) peak data rate with fallback modes of 5 Mbps, 1.6 Mbps, and 0.8 Mbps that are used to extend the range or to improve performance in high-inter-

FIGURE 2

Services Consumers Want from a "Telecom Bundle"

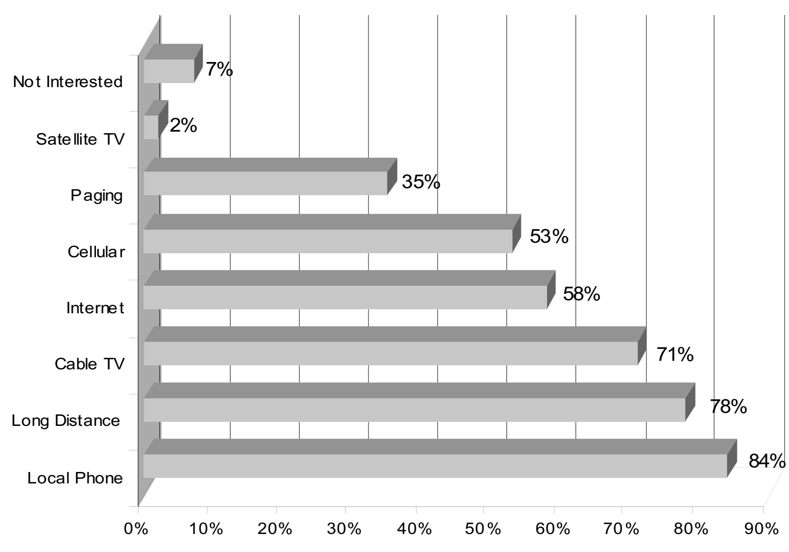

ference environments. It has powerful and effective security measures against eavesdropping, service denial, and unauthorized access, as well as active interference avoidance and mitigation techniques.

A key feature is its ability to prioritize up to eight simultaneous streaming-media sessions for audio and video. And it is unique in its ability to mix in four simultaneous toll-quality two-way cordless voice connections based on DECT.

HomeRF Capabilities Roadmap

HomeRF started in 2000 with shipments of 1.6 Mbps, data-only devices, and it was quite successful. Products from a variety of vendors were very easy to install and in fact were the first products used in the universal serial bus (USB), a tremendous advance.

This year the real focus has been bringing out 10 Mbps devices that handle voice, audio, and basic video, including cordless phones and simple streaming devices such as gateways, music devices, and Web tablets.

For the future, the Federal Communications Commission (FCC) is presently changing the rules in the United States to make it much more similar to the rest of the world in the 2.4 GHz band, and it is actually possible in the long run to scale this technology up to even 100 Mbps. But HomeRF has always been application driven, so as it moves forward it will look at additional video-streaming applications such as video tablets and set-top boxes and at data rates that will support them appropriately.

HomeRF Network Topology

HomeRF has a flexible network topology (see *Figure 3*). It is based on an integrated access device (IAD) at the core of the home network, providing a control point or referee functionality. The client devices can be quite varied, from traditional personal computer (PC) devices called asynchronous nodes (A–nodes) to appliances such as a Web pad that might do Web browsing to streaming applications such as games or videos. There could also be a very simple light client such as a portable Moving Pictures Expert Group (MP)3 player that plays directly off a DSL broadband connection or a traditional cordless-phone handset. In addition, while information flows predominantly from the IAD or central controller, all of these devices can communicate directly with each other. They do not need to loop back through the host.

Network-Layer View of HomeRF

An abbreviated version of the open systems interconnection (OSI) seven-layer network stack is shown in *Figure 4*. HomeRF, like most of these wireless networking technologies, tends to focus on the two lowest layers, the data link control and the physical layer (PHY) of the stack.

Most of the 802.11 systems were essentially wireless Ethernet, as were all of the legacy proprietary wireless data systems such as OpenAir and WaveLAN. They went up through the Internet protocol (IP) stack using carrier sense multiple access with collision avoidance (CSMA/CA). Similar to Ethernet, CSMA/CA is a random-access system in which nodes that want to access the network will randomly jump on the media after counting "empty" time slots since last known network activity. If by unfortunate circumstances two get on at the same time, they collide, back off, and try again later. The PHY is of course related to the 2.4 GHz band access and the appropriate modulations that are used.

HomeRF is a little different from other wireless networking technologies in the way that it works in this stack. The

FIGURE 3

HomeRF Network Topology

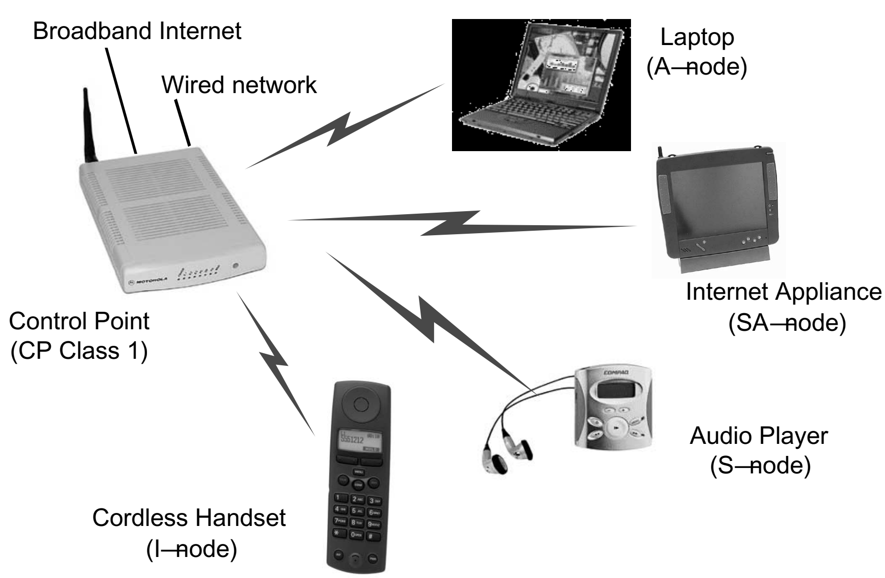

FIGURE 4

Network-Layer View of HomeRF

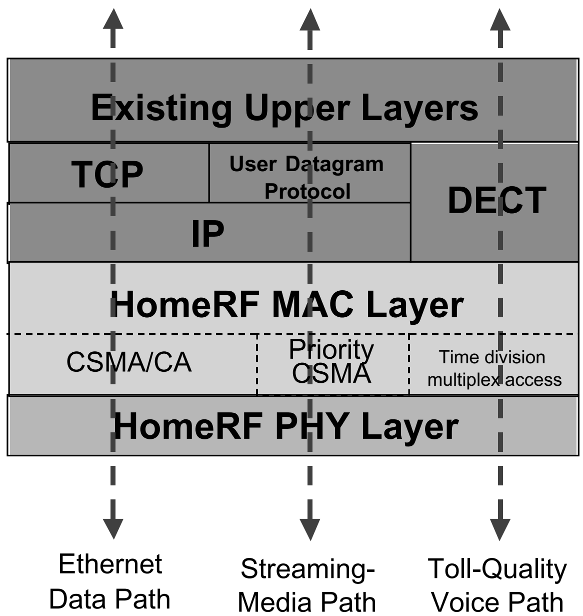

unique ability that HomeRF adds to the CSMA portion is the ability to grant some nodes higher-priority access to the media than others. Rules and regulations are present to prevent every node from being a high priority node and thus consuming all the bandwidth.

Another feature that distinguishes HomeRF is the ability to use time division multiple access (TDMA) as well as CSMA in the same frame. TDMA makes it possible to exactly, although dynamically, allocate slots for voice conversation to an individual handset. The handset is not precluded from going to voice over Internet protocol (VoIP) up an IP stack.

HomeRF: Media Access Control Layer Basics

Figure 5 shows the events that happen in HomeRF at the media access control (MAC) layer. The bulk of the time in a given frame, which is typically 10 or 20 milliseconds (ms) long, is dedicated to data networking, as long as there are no voice calls or priority streams. Within the data-networking time, streaming-media sessions get priority access. The time periods reserved for voice calls are based on the number of active voice calls. Voice calls will be divided up into multiple smaller slots depending on the number of calls. By reserving that time hard, there is no chance that data networking, printing a file, or Web browsing will interfere with the voice calls. That is a key attribute of HomeRF.

Dynamic Bandwidth Allocation

Another interesting and unique feature of HomeRF is that the proportion of voice calls to data networking will change depending on the number of calls. If there is only one call, only a very small amount of the available time will be used for voice calls. But four calls would actually require about half of all the available time in a given frame, so half the time would be dedicated to voice. It is all adjusted dynamically. So, for example, if a user has three handsets in his house and VoDSL service with three unique lines, it is not necessary to reserve three-eighths of his or her available bandwidth at all times just in case those lines ring. The system will simply use the bandwidth dynamically when it needs it.

Immediately Retransmits Voice Packets

Another unique feature of HomeRF is used if there is an interference problem, for example from a microwave oven. Less than 1 percent of homes in America have wireless networks, but 98 percent have microwave ovens. And wireless networks have to share the bandwidth with microwave ovens; thus the ovens cause interference. When that occurs, as soon as HomeRF hops to the next frequency, it will reserve time and immediately retransmit those voice packets ahead of any of the data networking events and prioritized streams.

Because the length of the frame is 10 ms, it can do all of that in a bounded time period. Thus it guarantees that in a given 10 ms period, there will be at least two chances to get data through. And hopping to a different frequency at a different spot in the band unrelated to the last is a very powerful way to avoid interference with voice packets.

The data does this as well, but on a much longer time period. The voice traffic has a pre-emptive prioritized capability to do it. This effectively squares the data rate. Since the two frequencies are independent, a raw packet-failure rate of a few percent (typical around microwave ovens) gets squared to much less than 1 percent—this means excellent voice quality with the chosen coder/decoder (CODEC).

HomeRF/Bluetooth PHY Layer Basics and Commonality

Band and Channel Access

The key thing to emphasize about the PHY is how it drives cost. It is a 2.4 GHz frequency-hopping system that is basically identical to DECT. In fact, both HomeRF and Bluetooth are basically identical to DECT. There is really no difference among HomeRF, Bluetooth, and DECT that has any substantial impact for the RF engineer. 802.11b is a different animal, however.

Modulation and Transmitter

HomeRF and Bluetooth use a very simple modulation, frequency-shift keying (FSK), and they use direct conversion and constant envelope modulation, which remove the need for linear power amplifiers and save a lot of complexity. Nominally they achieve 100 milliwatts (MW) saturated output power.

Transceiver Architecture

All of these wireless systems are half duplex. If they must be turned around very quickly, that adds cost. In the case of HomeRF and Bluetooth, there are 200 microseconds—an eternity—to turn around these transceivers using standard synthesizers on a single voltage-controlled oscillator (VCO).

Single-Chip PHY Complementary Metal Oxide Semiconductor Radio Frequency Integrated Circuit

Single-chip PHY complementary metal oxide semiconductor (CMOS) radio frequency integrated circuit (RFIC) is

FIGURE 5

HomeRF—MAC Layer Basics

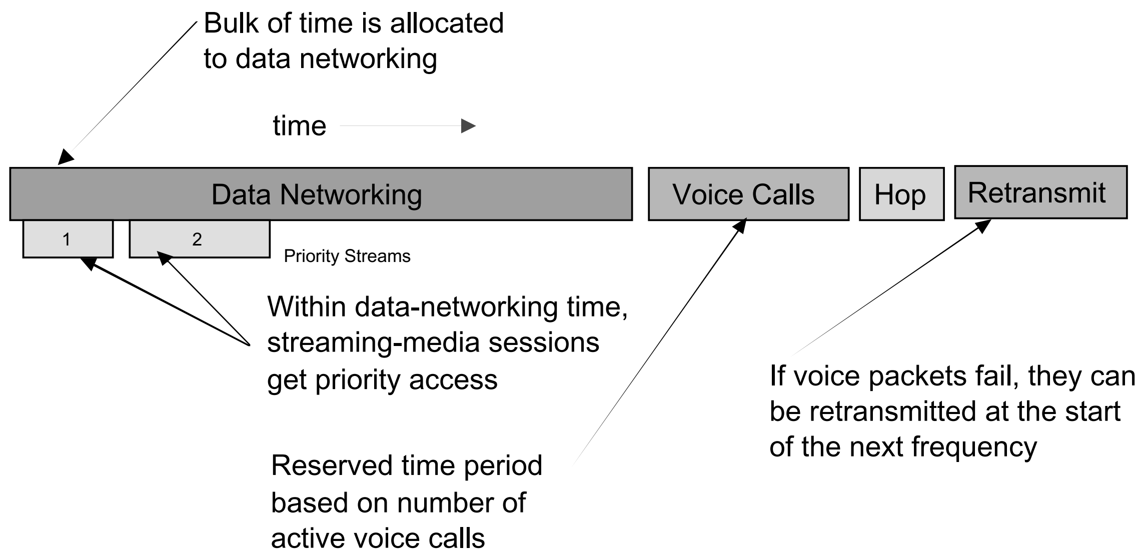

practical today for either HomeRF or Bluetooth, or dual mode for both. When people talk about the $5 radio, they are really talking about doing the entire radio on a single CMOS chip. This is already available for DECT and will soon be available for Bluetooth. The same technology is equally applicable to HomeRF.

Key Comparison Attributes between HomeRF and IEEE802.11b

The following sections will compare HomeRF with 802.11b, which is probably the most talked-about candidate right now outside of HomeRF or wireless home networking. The two are differentiated by their voice support, streaming media support, interference immunity, scalability, and security.

Voice Support

Voice support is an important feature. Consumer voice revenues in America are more than 10 times greater than the cumulative consumer/residential data revenues. And although the wireless local-area network (LAN) industry is very proud of having passed its first million-unit year, there are 50 million cordless phones sold every year in America. The cordless phone is a device that consumers are more familiar with, and it outsells wireless data devices by more than 50 times. The level of voice support is a major difference between HomeRF and 802.11b.

802.11b Has No Explicit Voice Support

Voice support is probably the single biggest hole for 802.11b—it has no prioritization for voice, no bounded-latency mechanism for voice, and no standard software stacks for voice features. And it cannot be Band-Aided to get voice support. The basic frame structure of 802.11b does not lend itself to any of the conventional voice-networking stacks.

HomeRF Was Designed to Support High-Quality, Multiline Cordless Telephony

HomeRF is the exact opposite. It was designed to support high-quality, multiline cordless telephony. The fact that it can map directly into DECT and use all of the service-provider CLASS features means that any cordless-telephone solution out there today can use all those standard networking features as it adapts to HomeRF. HomeRF handles up to eight simultaneous voice connections and has a 10 ms bounded latency. Frequency diversity and hopset adaptation provide high quality even with severe interference. And it is the only ratified global standard for multiline cordless telephony.

Streaming Media Support

Streaming-media support is another point of differentiation between 802.11b and HomeRF. Internet audio and video are key applications for broadband, making streaming media support important.

802.11b Has No Explicit Streaming Media Support

802.11b does not currently support streaming media overtly. Demos work fine, but network loading can wreak havoc on VoIP phones or streaming videos. 802.11b can handle a little low-rate MP4 video, but there are big problems with printing a file, placing a lot of networks in an apartment building, or bringing an interfering device nearby. Some of the worst problems are with cordless phones. But interference and scalability issues notwithstanding, 802.11e should eventually address much of this issue.

HomeRF Was Designed to Support Multiple Prioritized Streaming-Media Sessions

With HomeRF, self-interference is not an issue because it can prioritize. HomeRF can handle up to eight consecutively prioritized levels and either one-way (audio/video) or two-way (VoIP/videophone) traffic. Frequency diversity and flexible retry buffer lengths provide high quality even with severe interference.

Interference Immunity and Scalability

Sources of interference include Bluetooth, cordless phones, microwave ovens, and high-density housing with many adjacent networks. Interference immunity is very important, because when people pay for a service, whether they pay $5 or $30 a month, they expect the service to work. Consumers will not understand that because their Bluetooth device came home or their neighbor turned on a microwave oven, they cannot dial their sister or access their voice mail. That will be unacceptable.

802.11b Has Only Three Wide, Static Channels

802.11b divides the 2.4–GHz band into three wide, static channels. 802.11b always shares bandwidth with foreign networks. The channel bandwidth has a fixed location and is about 17 MHz. The whole network sits there without any mechanism for moving around. The only interference-avoidance mechanism is time diversity. If there is an interferer, such as a microwave oven or a cordless phone coming on, the network must wait for the interferer to leave—that is, for the microwave oven or cordless phone to turn off.

The problem is not so acute for data. Microwave ovens tend to be on only half the time on a 60-cycle repeat, so data can usually get through, although with sharply reduced throughput. If consumers are simply accessing the Internet, they are used to the Internet being flaky and slow anyway, and they blame any delays on the Internet rather than on the microwave oven or wireless products.

But the inability to avoid interference is obviously unacceptable when it comes to voice. The microwave oven shutting down the VoIP connection will be very unpopular. And a devastating problem can be caused by a fast hopper such as Bluetooth. In some respects it is worse than the microwave oven because it hops so fast. At 1,600 hops per second, there will be a Bluetooth transmission sitting in that one-third of the band once every 2 ms. That does not provide much opportunity to get the packets through, especially when it is only randomly accessing the media to begin with. Bluetooth is a very hard device on 802.11b.

Another problem is hit or miss, but when it hits, it can be much worse—namely, a variety of proprietary direct sequence spread spectrum (DSSS) cordless phones. Most of them are six channel, meaning they divide the band into sixths. And within the one-sixth that they randomly pick, once they close their link, they have their call. As long as they have an acceptable packet-error rate, they remain static and will just sit there. For example, if one of these cordless phones closes a link and happens to be on the same channel as an 802.11b, the phone actually has something like a pilot carrier and it also operates full duplex. And it will just sit there continuously and shut the entire 802.11b network down—not to low throughput but to no throughput. The network is dead once that happens.

HomeRF Has Narrower, Dynamic Channels

HomeRF, of course, is very different for two reasons. First, the channel bandwidth is much lower, at 1 or 3.5 MHz. Thus rather than using up 1/3 of the band at a time, HomeRF uses in the voice mode only 1/75 of the band at a given time. On top of that, when interference does occur, HomeRF is always in motion, moving around to different parts of the band. That allows it to use both frequency diversity and time diversity to get around interference.

And HomeRF has the ability to adapt its frequency hopping on the fly. So if it encounters interference in a certain portion of the band, it can try to adapt around it so that it spends less time in that band, or at least never spends time in that band two times in a row, making the retry mechanism is most effective.

HomeRF also ignores foreign networks for graceful degradation in high-network-density environments.

Security

Regarding security, PHY/MAC attributes are critical because end-to-end solutions are generally not available in the home. Security is an important issue, although it can also be greatly overblown. Still, the issues that consumers care about are important ones for the industry, even if they are not always rational.

802.11b Has Three Well-Publicized Security Flaws

802.11b has three security flaws that have been highlighted both by the popular press and by industry experts. The first flaw, weak encryption with only a poorly defined, 24-bit initialization vector (IV) management, is completely overblown. It is unfair to 802.11b because even if MAC layer data was decrypted, it is usually useless without full access to the applications associated with the data.

The second flaw, however, is real. Open network access says that with an 802.11b device, a user can troll around the homes in his neighborhood looking for an 802.11b gateway. It is pretty trivial for the uninvited user to get on. Once on, the user could get a neighbor's DSL service that he or she was not paying for, or even get onto a neighbor's network, check out his or her hard drive, and look at his or her phone logs.

A third problem that has not been publicized as well is mass denial-of-service susceptibility—basically, a protocol attack. A good example of this is when people shut down the CIA Web site or MSN or Yahoo. This is vandalism. It is the electronic equivalent of painting on subway cars and in 802.11b, unfortunately, the protocol works such that the control signals are in the clear and are not network–ID dependent. It is very simple to build a device that goes into a neighborhood and broadcasts instructions to shut down all the 802.11b networks.

HomeRF Addresses All Three Concerns

HomeRF thought about these issues and addressed them. The solution for the first concern was to add 128-bit encryption and develop a tamper-resistant 32-bit initialization vector.

But the big difference, and this one is much harder for 802.11b to deal with, is that in HomeRF there really is no equivalent of the 802.11b open-access mode. Furthermore, a spec-compliant device has no way to pass what are called "promiscuous packets" up above the MAC. "Promiscuous packets" means listening for any packets that are valid HomeRF packets but are not on the listener's network ID. In

HomeRF, a spec-compliant device can not pass those up. In 802.11b, because it uses upper-layer techniques for authentication and access, it basically is required to pass them up. And that makes drive-by hacking very easy to do. With HomeRF there simply is no equivalent to this.

With HomeRF, mass denial of service is virtually impossible because of its frequency-hopping PHY and ignoring of foreign network IDs. If there are three static channels and a low-level protocol command to shut them down as in 802.11b, then it is fairly easy to jump across three channels that are fixed in time and shut them all down. With HomeRF however, multiple networks in a neighborhood or apartment building are all on random hopping sequences that are not known in advance. They are on different phases and time bases. They also ignore anything that does not come from their own network ID. And they have a unique 24-bit network ID that is not known in advance. A device that can figure out these network IDs, time bases, and phases and send low-level protocol commands could be built by the CIA, but not by drive-by hackers.

Other Comparison Attributes

Cost

In the end, cost will be the most important issue. HomeRF starts from a cordless-telephone cost basis, but the market sets the price. HomeRF is basically the same as DECT from a cost point of view. Thus it is cheaper than 802.11b.

But cost is one thing; price is another. Some very aggressive pricing among the 802.11b services has obscured HomeRF's cost advantage, and thus cost has not yet been as big a differentiating factor for HomeRF as it will be in the long run.

Range

HomeRF2.0 and 802.11b have roughly similar ranges.

Peak Throughput

HomeRF 2.0 and 802.11b have roughly similar peak data-only throughput. HomeRF actually has much more significant throughput if there is a lot of interference present, but outside of that case they are similar.

Power Consumption

Power consumption is a big differentiator for HomeRF, just as it is for Bluetooth. And HomeRF can be done with much, much less power than 802.11b. HomeRF power-saving devices consume less than 10 milliwatts standby with full transmission control protocol/Internet protocol (TCP/IP) connectivity.

Network Topology

HomeRF uniquely supports host-client and peer-to-peer simultaneously.

What about IEEE802.11a?

More and Cleaner Bandwidth

802.11a products operate in the 5 GHz band, where there is more and cleaner bandwidth. There is far less interference because that band has 300 megahertz (MHz) total in the United States today, and probably closer to 470 MHz total in Europe, although only 100 MHz so far in Japan. Unfortunately, it is different in all those countries. By contrast, a beautiful thing about the 2.4 GHz band is that it truly is global. The rules are close to identical globally, and the band allocation is virtually identical globally. That remains a big advantage for the 2.4 GHz band over the 5 GHz band.

Same Ethernet-Oriented MAC Layer as 802.11b

802.11a is basically the same as 802.11b in the sense that it is wireless Ethernet. It still lacks explicit voice and streaming-media support, relies on upper layers for security, and has the same security issues as 802.11b. But 802.11a can put 54 Mbps in 12 channels at 5 GHz, while 802.11b puts 11 Mbps in three channels. That is a 10-fold increase in system capacity.

Conclusion: Ideal Solution for the Enterprise

Overall, the far greater system capacity of 802.11a makes it an ideal solution for very dense enterprise deployment. For offices of the future, or even offices of today where they do not put in wire, the system capacity that 802.11a represents is tremendous.

Market Positioning

Will this technology come home? Some of the core technologies used in 802.11a will be very applicable to the HomeRF roadmap, and 802.11a will likely come home just as 802.11b has. There is no question that these will show up in people's homes. But the center of the home market is really HomeRF because of its cost, security, toll-quality voice, and streaming media.

Nothing can provide ad hoc connectivity and low power like Bluetooth, and future mobile markets will be Bluetooth-centric. It is a tremendous technology that also is low cost.

In the business market, 802.11a is where the future is going to be, because it offers wireless Ethernet, performance, and roaming.

Summary

HomeRF combines 10 Mbps data, toll-quality voice, and streaming media for the broadband Internet home. It is uniquely optimized for the broadband Internet home, and it can be scaled to more than 50 Mbps data rates under new FCC rules.

Addressing the Deployment Dilemma

How to Select the Right Remote Terminal Enclosure for Sophisticated Broadband Equipment and Enhance its Revenue-Generating Capability

Minesh Patel
Formerly Director, Product Management
Marconi Communications

Deploying outside the plant to bring high-revenue–generating broadband services to residential subscribers finds service providers trapped between two fundamental issues: technical requirements of the sophisticated broadband equipment and right-of-way restrictions imposed by local governments.

On the one hand, broadband equipment consumes more power, requires more battery backup, generates more heat and, in some cases, is more complicated to connect than conventional telephone equipment. This suggests larger, more robust enclosures and the potential for greater acoustical noise from fans and air conditioners.

On the other hand, municipalities and other local governing bodies, responding to complaints of property owners, are increasingly reluctant to allow deployment of additional, larger, noisier enclosures in residential neighborhoods.

This is a challenge to service providers and enclosure manufacturers alike. The provider must understand the physical requirements and limitations of broadband equipment. And the manufacturer must be able to develop an "enclosure system" to contain the equipment, maintain the right thermal environment, allow room for ample battery backup, provide physical access for service, and maintain form factors and footprints compatible with local restrictions at a reasonable cost.

Service Providers' Responsibilities

Enclosure manufacturers are already addressing "broadband readiness" issues so suppliers can obtain new enclosures to accommodate the new services. However, providers need to know what to ask and what to tell the manufacturer to get the proper solution. And they need to understand what manufacturers claims and specifications actually mean—particularly new competitive local-exchange carriers (CLECs) that may have limited understanding of the technical issues and choose to rely totally on the enclosure manufacturer's expertise.

A manufacturer might say, "This cabinet cools 5,000 watts." While it may cool 5,000 watts when the cabinet is empty, the ventilating system may not be able to cool anywhere near 5,000 watts when equipment is installed. The reason is "system impedance." Equipment inside the cabinet will restrict airflow and cause a pressure drop that reduces the amount of air moving through the electronics and heat exchangers—thus reducing the amount of heat that gets removed.

Or the manufacturer may recommend the addition of a fan shelf to increase cooling. Depending on where it is placed and how it is operated, the shelf can, in fact, *reduce* the cooling. The fan shelf may only be operating at 200 cfm, while the cabinet thermal system may be operating at 600 cfm. In this case, the slower moving fans in the fan shelf, instead of helping, have further impeded the airflow and therefore diminished the cooling effect.

The service provider must tell the cabinet manufacturer what specific equipment the cabinet will contain, and the manufacturer must have the skill sets to know that equipment and how it will affect the thermal system within their cabinet. The equipment layout and relative positioning can also be a key to successful "packaging." Service providers must be flexible to allow the configuration to fit the enclosure as much as the enclosure fits the configuration.

Thermal Management

Proper thermal management is necessary to ensure the dependable operation of sophisticated heat-generating equipment that requires the protection of a sealed enclosure. The heat must be transferred from the points at which it is created to the outside air. In a broadband-ready enclosure, this becomes a complex problem, as illustrated in *Figure 1*.

Heat is generated at the silicon die level and transfers through the chip to the component, the card, the sub-rack, the rack, the air circulating inside the enclosure, and finally to the heat exchanger. Since the broadband enclosure may contain hundreds of these heat sources, with many components, cards, racks, and other obstructions (system impedance) to air flow, it becomes a very complex problem for the enclosure manufacturer to remove all of this heat in a cost-effective manner.

FIGURE 1

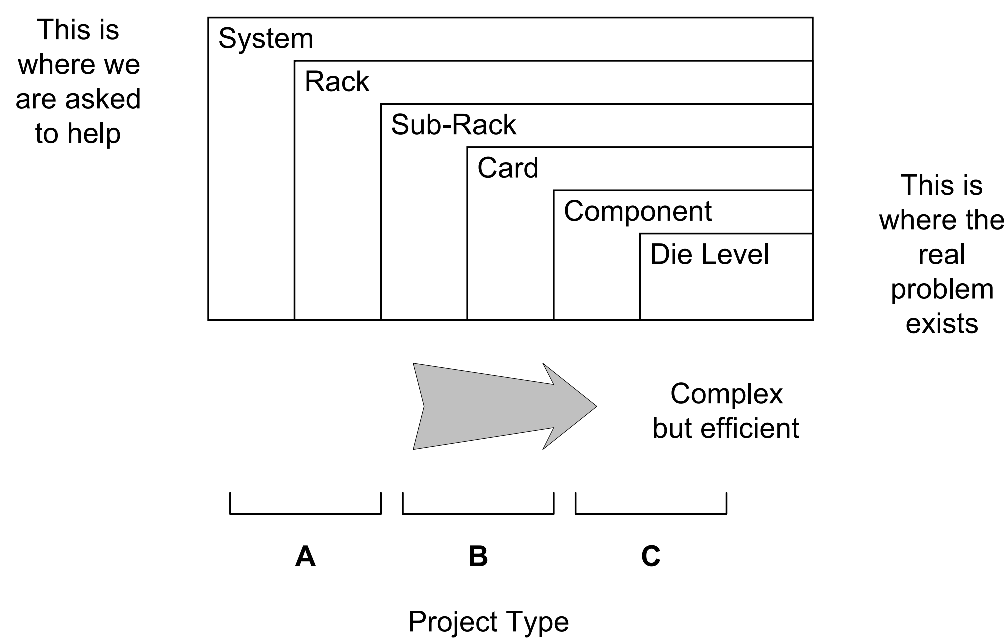

The Telecordia GR-487 (General Requirement) provides specifications for a sealed heat exchanger system that the telecommunications enclosure industry follows. While most incumbent local-exchange carriers (ILECs) are familiar with Telecordia specifications, newer providers may not be, or may overlook the fact that the specification may not have been followed. Some may be unwilling to pay for the cost of a properly engineered sealed system, since low initial cost is important to their business plan.

However, when equipment fails or need for expansion occurs, the cost to upgrade is considerably greater than installing properly designed equipment with room for expansion in the first place.

Thermal "redundancy" of the enclosure thermal management system—its capacity to handle "unforeseen" thermal events—is a critical concern when selecting an enclosure. With all of the heat generated by broadband equipment, a catastrophic failure can leave very little time before temperatures exceed 85°C, the upper limit for most electronic circuitry.

Air-quality requirements within a broadband-ready enclosure are the same as those for previous generations of cabinets containing sensitive equipment. The objective is to keep humidity and contaminants out as much as possible with a sealed enclosure design that employs a heat exchanger system.

Thermal Control and Acoustical Noise

Acoustical noise is an important concern when using high-speed fans and motors. Many remote enclosures are going to be placed in residential neighborhoods (sometimes close to bedroom windows!). To move the amount of air necessary to cool broadband equipment, fans are going to be pretty noisy. How do you minimize that noise and still maintain the desired thermal environment? With a thermostat controlling the whole system, everything turns on whenever the contacts close. However, with a more sophisticated controller that monitors or "polls" the environment, you may determine that only the fan shelves need to be on at a given moment, but not the heat exchanger fans—thus reducing the amount of noise.

Thermal monitoring is also required to assure that, if the cabinet temperature exceeds 65°C, appropriate alarms (including fan failure alarms) are sent to the network alarm center. Usually, there are separate alarms for environment and power distribution (failure) conditions; but from a "system" perspective, it is generally more economical to combine all alarms in a single package that communicates with the network alarm center.

"Scalable" Cooling

In some deployments, you may not need a full broadband cooling package and everything that goes with it. Providers should realize that cooling can be scaled down and cost taken out of the enclosure. If you give your enclosure manufacturer a complete description of the type and amount of equipment to be housed, the manufacturer can size the fans or heat exchangers appropriately. However, this must be done in the factory. If you find you need greater thermal capacity later, it will be difficult to "scale up" without interrupting service to your customers. In any case, the field is not the place to tamper with the capacity of critical cooling systems.

Battery Backup and Cabinet Space

With broadband, the equipment consumes more power and, to get the customary eight hours of battery backup, you need more battery space. In many older cabinets, slide-out battery drawers held top-post batteries. The drawers and support brackets consumed a lot of volume, making battery compartment utilization poor at best.

Today, new front-access batteries with a different form factor are available and drawers are being eliminated. If the provider has standardized on a particular battery, the form factor of the battery compartment in the cabinet they choose must be compatible with their battery and allow the most efficient utilization of that space inside that cabinet. If the utilization is low, backup capacity will be reduced. The choice becomes "choose another battery" or "choose another cabinet," depending where the provider is in the enclosure selection process.

As with thermal systems, the provider can save money on battery backup by knowing what type of equipment a particular enclosure will house. If it is high-speed data equipment and not lifeline telephone service, the provider should decide if a less-than-eight-hour battery backup system is adequate. The cabinet manufacturer can then configure and equip the cabinet accordingly.

A Cabinet As a System

Often, the enclosure is viewed simply as a protective "wrapper" around the revenue-generating equipment. But today, you can't always get approval to increase the size of the "wrapper" to hold the additional equipment required by a broadband terminal. So you need a broadband enclosure with the flexibility to integrate as many of the components as possible within the enclosure—in other words, a "system."

Enclosure "system" components include the power-conversion equipment, power distribution, environmental monitoring and control, battery management, and cable management. All of these components take up space. And with broadband deployments, you need "more" of each component. But "more" of these components limits the space for the revenue-generating electronics.

The system concept can lead you out of this dilemma. By requiring manufacturers to design broadband enclosures from a system perspective, network providers can stimulate manufacturers to be more creative and find ways to pack more into that wrapper—but still give you the "craft friendliness" that you want.

For example, with the increased cooling requirements, larger fans will be needed to drive air through denser heat exchangers. Unfortunately, the power curve for these fans is exponential—doubling the speed of the fan may only increases the airflow by 25 percent. Rather than fuse these larger fans with a separate distribution panel, the power distribution can be integrated within the same device that provides the environmental monitoring and control.

As another example, when you deploy a remote terminal, you normally deploy a separate standalone transfer switch and a cross-connect at that same location. Integrating a power transfer switch as part of the cabinet system is another way to cope with the right of way issue—and eliminate another headache for the provider.

The more thought you and the manufacturer put into your cabinet before it leaves the factory, the more suitable it becomes—as a system—for your broadband application.

Fiber and Hybrid Management

With fiber, copper, and coax entering and leaving the remote broadband enclosure, media management and protection become critical issues in enclosure design. There must be enough room to attach the fiber troughing that maintains minimum bend radii and appropriate spacing—and enough room for the equipment. There must also be enough space between, above, and below the equipment for fiber management.

If an enclosure is poorly designed, it may have enough room to hold the equipment at initial deployment; but when maintenance or troubleshooting are required, there's no room to get your hands in to handle the cables to do connectorization, routing, lacing, etc. There are many systems for fiber splicing and patch panels, and providers have their own preferences. It is the cabinet manufacturer's job to provide adequate room for the components the provider selects.

The service provider must know where the access ports are on the chosen equipment when talking to the cabinet manufacturer. Location of the ports varies greatly. It may be front access or rear access. If it's rear access, do you need a swing frame? Or can you get to it some other way? Swing frames are costly, heavy, and take up cabinet space. Before settling on swing frames, see if there is a way—via configuration—that your enclosure supplier can save you from the additional expense of swing racks.

Ease of Deployment and Maintenance

Now that you have much more revenue being generated out of your cabinets, you can't afford excessive downtime. Nor can you afford to disappoint the many customers who may be served by a single cabinet. How easily can you install your enclosure? After you've installed the cabinet, how easy is it to get the cables in and out of the cabinet in a safe manner? With higher-heat equipment, if something does go wrong, you have to get to the site quickly, and equipment must be quickly serviceable when you get there.

The provider should ask, "How much maintenance am I going to have to do? Am I going to have to go out often to clean filters, replace fans, and replace heat exchanges, etc.? Or do I just deploy the enclosure, let it go, and only service it every four or five years when the fan burns out? And when something does need replacement, how easy is it for me to replace? Will it take eight hours to replace a fan? Will I have to tear equipment out of the cabinet?" All these questions should be asked before you choose an enclosure—not after it's deployed in the field.

Compartmentalization

Compartmentalization is an extremely important issue because different craft grades will be accessing the cabinet. With sensitive electronics, providers often want to restrict access to different parts of the cabinet, and enclosure manufacturers have stratified the crafts by compartmentalizing the enclosures and installing different padlocks on each compartment.

Not only does this restrict crafts to the appropriate compartments, but when a technician is working in one compartment, only that compartment is exposed to the elements. Splicing could be going on eight hours a day, for a couple of days, with the doors wide open. Without compartmentalization, the equipment would be exposed to heat, dirt, moisture, etc. With it, equipment in other compartments remains undisturbed in separate controlled environments.

Absorbing the Cost of Broadband Readiness

While new broadband enclosures may be restricted to present form factors and footprints by municipal mandates—and may look much the same outside—do not be confused by the similarities and do not confuse the costs. There is an added cost for cooling, for battery backup, and for media management. But when you and your enclosure manufacturer work together, you will find creative ways to absorb the added costs and enjoy the increased stream of revenue that your broadband-ready cabinets will help you generate.

Scaling to Handle DSL Demand

Mark Peden

Vice President, Technology Marketing

Simpler Networks

This paper discusses scalability, the number-one issue network service providers (NSP) face today as they deploy digital subscriber lines (DSL). It presents the issues that prevent the industry from achieving scale. The paper begins with a brief introduction to loop management and then discusses recent advancement in loop qualification, provisioning, testing, line maintenance, and, finally, service assurance.

Yesterday's Provisioning Process

Figure 1 shows yesterday's provisioning, an expensive and time-consuming process. It was very manual, laborious, and error prone. The end user placed the order either over the Web or with a phone call. It involved some order loop prequalification that, in the case of a competitive local-exchange carrier (CLEC) or data local-exchange carrier (DLEC), could take from a week to 10 days.

The second step was to get a local service request (LSR). Again, the CLEC or DLEC itself actually ordered the service or the pair, depending on the type of service, from the incumbent local-exchange carrier (ILEC). The next step, scheduling the workforce, took anywhere from a month to a month and a half. Designing and assigning the digital subscriber line access multiplexer connection point (DSLAM CP), which included setting up the virtual path identifier/virtual circuit ID (VPI/VCI) modeling, took about another week. Manually provisioning DSLAM took another one to three days. Rolling a truck took perhaps two or three days. Then the bills started to come.

In the early days, even the well connected needed to beg and plead with the local telephone company to get DSL.[1] The reason? The digital loop carrier (DLC) was between the requestor and the central office (CO). After finally agreeing to supply DSL, the telephone company said, "Look, we are going to get it for you, but you cannot tell anyone." That is because the telephone company had to re-engineer the loop plant to deliver DSL service. But others could not get this exciting, wonderful service for another two to three years.

The telephone company hangs notices on doorknobs and slips ads in every bill to promote cable service. "Call us today—no contract necessary, only $19.95 a month. Within one week, we will install it and guarantee that it works." These marketing efforts reflect how far cable service has come; real progress has been made. No service provider anywhere is promoting DSL like that, so DSL has not reached the mass market.

Reaching Mass Market

What does it mean to reach mass market? Mass market is not 1,000 lines a day; it is not 5,000 lines a day. It begins to happen when a single service provider can actually deploy 10,000 lines a day. No service provider in the industry today can actually do that.

How can the industry position itself to reach mass market? By meeting these important requirements:

- General product availability
- Standardized products
- Product interoperability
- Consumer awareness
- Product and service alternatives

Product Availability

DSL products must be generally available. The physical DSLAMs must be installed, or, better yet, the next generation of digital carriers deployed in COs or remote terminals.

Standardization and Interoperability

Products need to be standardized to allow standards-based interoperability. The DSL Forum has made significant progress toward facilitating standardization. At the last SUPERCOMM, more than 30 vendors demonstrated standards-based interoperability. It was exciting to see the whole industry come together and demonstrate standards-based products and how they could all work. Those efforts and work within the DSL Forum will bring the industry to the point where it can offer a retail model. Then anyone can go to a retail store, buy a product off the shelf, go home, plug it in, and be up and running fairly quickly. The industry is progressing toward that goal, but it is not there yet.

Consumer Awareness

Consumer awareness has increased. Everyone is talking about DSL now. People know what it is; they actually understand what it means and the value that it can provide. That is rewarding for the people who have worked on this technology for so long.

FIGURE 1

Yesterday's Provisioning Process

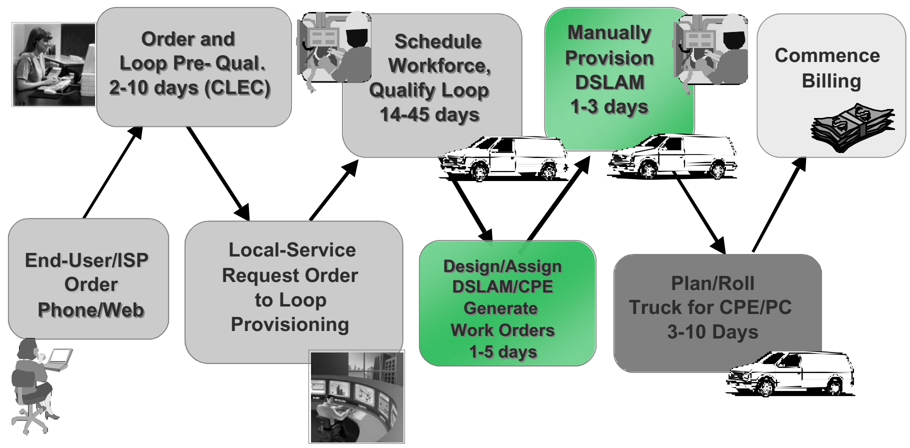

Product and Service Alternatives

Both ISPs and competitive access providers (CAP) have begun to offer some product and service alternatives that add value by rolling out services, addressing security, voice over, video on demand (VOD), voice over digital subscriber line (VoDSL), and virtual private networks (VPN). Koreans use the term KOD–karaoke on demand. Karaoke is very popular is Asia; it is part of the social scene. Now it is actually being delivered over DSL. Servers provide localized content caching and deliver karaoke directly to the end user. That is just one example of how creatively people are using and applying DSL.

Today's Provisioning

Provisioning still presents some challenges, such as

- Accurate loop provisioning
- Support for electronic bonding
- Establishment of physical connections
- Errors resulting from manual processes
- Correction of impairments

Accurate Loop Provisioning

The two key barriers to accurate loop qualification in the provisioning space are false positives and false negatives. A false positive occurs when a customer learns DSL is available and calls or goes on-line to order it. The service provider makes a quick assessment, perhaps some book qualification, and worse yet, some narrowband qualification. Believing it can deliver DSL, the service provider spends some money, tries to get a line to the customer, and perhaps delivers a new line. After spending time, money, energy, and personnel resources, the service provider discovers it cannot provide DSL to the customer. That very costly and very real issue affects about 20 to 25 percent of today's orders.

A false negative occurs when the service provider says a customer cannot get DSL. This customer operates over the border in a gray area. The service provider feels strongly confident of where it can and cannot deliver DSL. So it tells a customer that is just outside of its area, "No, we cannot get it to you." However, by spending a little money, the service provider could provide DSL to the customer.

In California, when Pac Bell first rolled out stored program control (SPC), it limited loop reach to 12,000 feet. Why? Because it knew that 99.99 percent of the time it could deliver services to customers within that range. Trials and tests of the technology concluded that it could deliver with a loop reach of 23,000 feet. And yet Pac Bell decided to limit it arbitrarily to 12,000 feet. Doing so made Pac Bell highly confident that it could have DSL up and running and deliver the service consistently.

Generally, some carriers initially limited the service and then expanded it over time. They placed some restrictions on loop reach because, with a shorter loop, they could be more confident about the delivery of service. Over time, some service providers extended the loop reach to service a wider customer base.

Electronic Bonding and Physical Connection

Electronic bonding is enabling true flow-through provisioning. Once the loop has been qualified, it occurs between Internet service providers (ISP), potentially between DLECs or CLECs, and then between ILECs. Another key issue is establishing the physical connections, especially in a competitive environment where providers might use a line-shared loop. The industry is still working through yet another key challenge that concerns splitter usage and positioning. Different models currently have industry support. But there is no one simple answer, because competitive carriers and their unique challenges make the issue more complex.

Errors and Impairments

Manual processes frequently result in errors, and impairments in the loop may need addressing by the provider once the loop is qualified to address data-rate options through testing. Many service providers automate their testing using the Class-5 switch. They connect the test head and qualify through the Class-5 switch. However, that is a form of a narrowband testing used for *broadband* qualification.

Using a narrowband test does give throughput options. It does no spectral testing per se, and spectral issues might potentially exist within the binder. A narrowband test does tell with some level of accuracy whether a customer can actually get service. Service providers consistently report that 75 or 80 percent of the time, a narrowband test accurately predicts whether they can actually deliver services.

A number of service providers are working aggressively to reach the mass market. As they approach the goal of 10,000 lines per day, their narrowband tests may result in as many as 2,000 false positives or false negatives. That means they cannot provide a service-level agreement (SLA) if they use a narrowband test, and they cannot guarantee data rate throughput.

Post-Qualification Concerns

Many concerns must be addressed after loop qualification. Interconnection issues between the ISP, CLEC, and ILEC become critical with throughput. So do issues related to the CPE–to–DSLAM auto configuration. The industry is working to develop a common model. Many service providers now support Tech 59, an approach being developed within the DSL Forum. Simplified installation with plug and play (PnP) is also widely supported. Both improve consumers' experiences with DSL. They can buy the product, take it home, plug it in, and use it fairly quickly. The news is good—many service providers are succeeding in providing a PnP model that consumers can install.

External Factors

Carriers are committed to and focused on providing long-distance services. They consider competitive access key to facilitating their entry into that market. To accomplish that, a number of external factors must be considered:

- Regulatory issues
- Emerging technologies
- IP and ATM issues
- New network and content methods
- Usage drivers

Regulatory Issues

Public utility commission (PUC) issues vary from state to state. In Illinois, for example, the state's PUC and the local carrier were in conflict about the former's ability to deploy in a DLC environment. The carrier decided to hold or suspend some deployment until hearings were held and the issue was resolved. State issues affect deployment, as do federal issues.

The Network Reliability Interoperability Council (NRIC) is an organization endorsed by the Federal Communications Commission (FCC). Its mission is to develop policy recommendations for the FCC. To some extent, the NRIC is a subset of Committee T1/E1.4. Many who have been involved in standards work within committee T1 are also involved in the NRIC. The council has posted some of its policy recommendations on the Web at www.nric.org.

Emerging Technologies

New technologies are constantly being developed and deployed. Some new technologies, like repeaters, are known disturbers of DSL deployment, but others are actually more spectrally friendly to different DSLs. For example, G.shdsl is more spectrally friendly to asymmetric digital subscriber lines (ADSL). Still other new technologies being deployed provide some interesting challenges in the loop.

IP, ATM, and New Technology Issues

Internet protocol (IP) over asynchronous transfer mode (ATM), or IpoA, is a different issue. New network content models are being developed for voice over digital subscriber line (VoDSL). Quality of service (QoS) is coming. So are usage drivers, new content, and applications for home networking drivers.

Ongoing Challenges

The industry must tackle availability constraints immediately by helping providers to complete CO rollouts and helping the next-generation digital loop carriers (NGDLC) provide ubiquitous service. SBC has earmarked $6 billion for buildout, much of which will go to NGDLC. SBC has taken the leadership and made a commitment to addressing the residential market.

Accurately qualifying and automating what remains of the manual provisioning process continue to be goals. The industry must provide support for QoS, service-level agreements (SLA), VoDSL, and some value-added services. As new technologies are introduced, they will encourage DSL deployment, increase the number of customers served, and enable new content, home networking, and other access devices.

Loop Management

The DSL Forum undertook a project on loop management systems (LMS). A white paper describing the LMS project, "DSL Form 2001.067," is now available at the DSL Forum Web site (www.dsl.forum.org). A white paper is also available on LMS.

The project summarized some problems described in this paper: increased operating complexity, slow response to customers, poor QoS assurance, managing the line-sharing environment, high customer-service churn, and poor loop record. Then it attempted to identify some solutions, including loop testing and qualification, better test-access functionality, service migration and provisioning, line sharing, and automating sub-loop unbundling.

Figure 2 is a diagram of the LMS architecture. The system contains multiple elements to automate testing, provisioning, and maintenance functions, and it enables service providers to install, troubleshoot, manage, and maintain existing and emerging services on the local loop more efficiently.

The LMS architecture includes network operations, operations support system (OSS) access, a control interface, a test head, test access, a metallic access matrix, and a splitter. The four basic items in the functionality are provisioning (physical switching), test access, optional test-head or devices, and remote management.

An LMS system could exist in a co-location cage where a competitive carrier may use it. *Figure 3* shows how it could exist within the CO just off the frame or at an intermediate distribution frame; some carriers might potentially use automated or remote switching within a DLC. Some

FIGURE 2

LMS Architecture

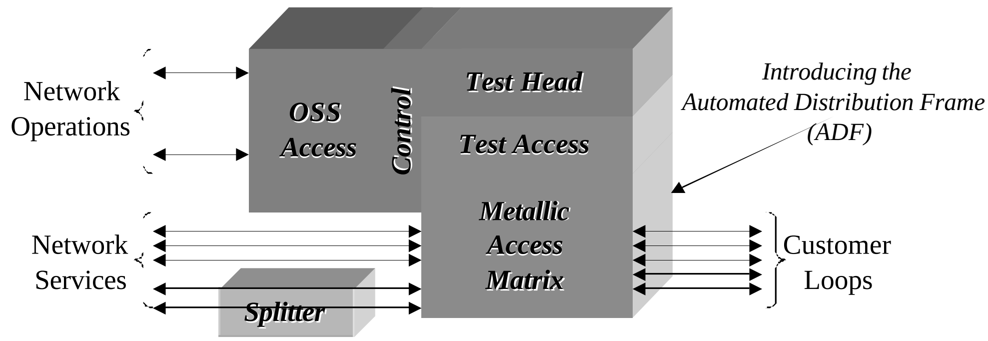

providers are looking at introducing DSL into commercial buildings—for example, in hotels or motels—where they can automate switching or introduce high-speed services at the remote terminal.

Current Needs

To develop an automated system, a relay is necessary. A relay makes a physical connection that enables any type of cross-connection. Typically a relay is on a relay board, and it is big. A few companies have been working on micro relays similar to the one illustrated in *Figure 4*. The micro relay is a solution to the space limitations of loop management.

The development of the micro relay is also an effort to address some of the industry's "future" needs; in particular, to provide real-time testing and qualification of access to copper lines and to make full integration possible. Testing can no longer be limited to narrowband, which, as mentioned earlier, is the most common method of broadband qualification because it can be automated through the classified switch. Efficiencies of space and power; optimal use of DSLAM or DSL ports; easy, cost-effective ways to upgrade customers; and the ability to address both protection switching and reconfiguration—all of these are important factors in providing technical solutions. In addition, some regulatory issues must be addressed through technical solutions.

Broadband Distribution Frame

Based at the CO, yesterday's deployment included provisioning at the frame, where someone physically connected pairs (see *Figure 5*). A new concept, the broadband distribution frame (BDF), helps to provide any-to-any, non-blocking, protocol- and vendor-independent distribution. *Figure 6* illustrates its support of flow-through provisioning.

By coupling inputs and outputs remotely, BDF automates many functions that required site visits. It automates testing, provisioning, and maintenance functions in supplying DSL services, activities that become even more complex in a remote terminal or a DLC environment. BDF can provide test access to any loop, and testing is no longer limited to narrowband testing. The tests can be full broadband or full spectrum. BDF also allows independent management of DSL and plain old telephone service (POTS) on the shared loop. It keeps configuration records and provides easy

FIGURE 3

LMS Applications

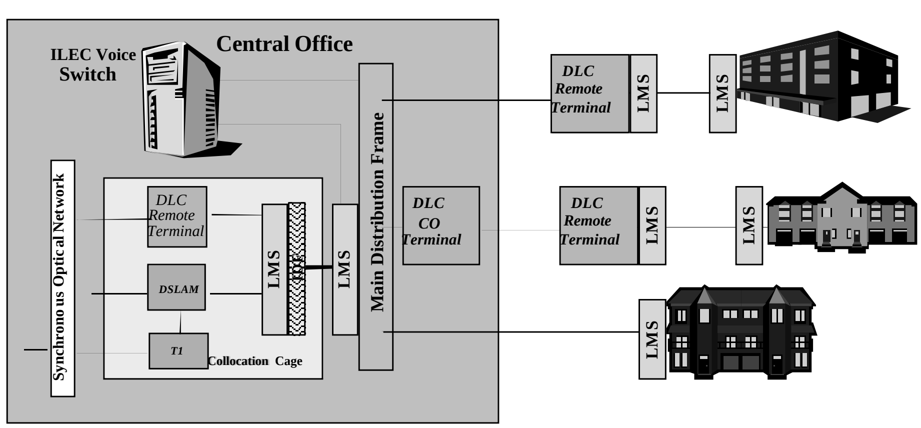

FIGURE 4

A Micro Relay

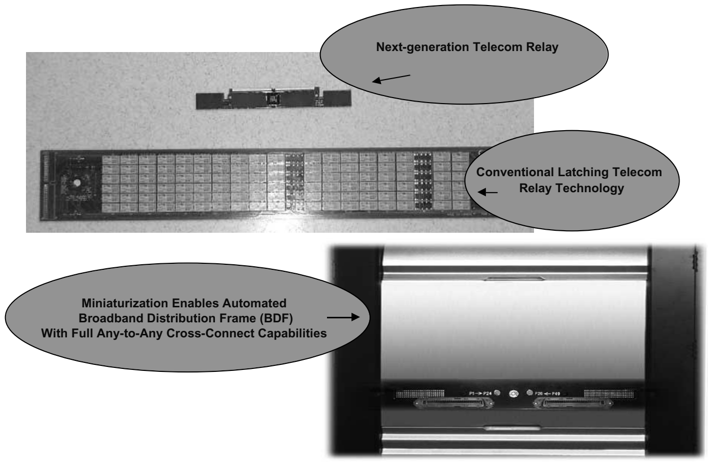

implementation of service changes and integrated protection switching.

A BDF or an LMS with automation can actually address and potentially eliminate the workforce qualification, which could take up to 10 days in the manual provisioning process. Finally, auto-configuration techniques now actually address the designing or assigning of the DSLAM customer-premises equipment (CPE) via auto configuration techniques. Automating the loop-management element can potentially reduce the waiting period from weeks to only a few days.

How BDF Works

Figure 7 is a graphical representation of how BDF works. Many companies offer fairly similar products in which all pairs basically come through a frame, perhaps in a remote environment. The CO in the corner could also be a remote terminal where the pairs simply pass through an additional box.

FIGURE 5

Yesterday's CO–Based DSL Deployment

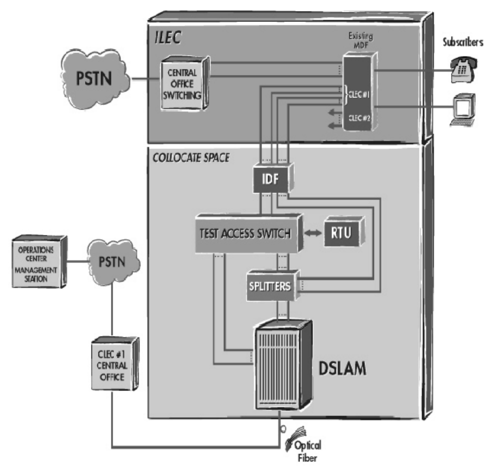

FIGURE 6

Flow-Through Provisioning with a BDF

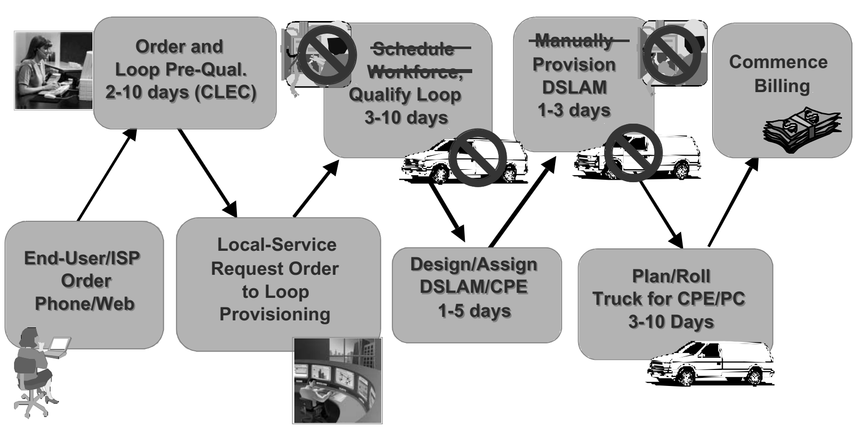

FIGURE 7

Pre-Wired BDF and DSLAM

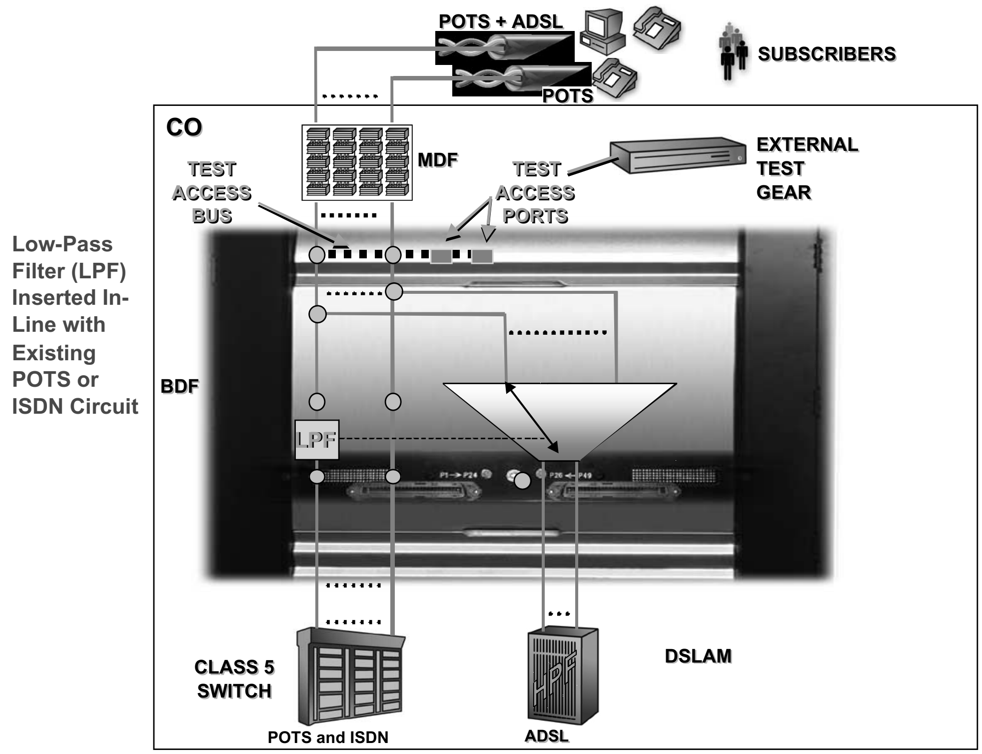

Some lines coming into this product perform what is considered a half tap and actively terminate a DSLAM. The low-pass filter (LPF) shown in *Figure 7* can be remotely inserted in-line with an existing POTS or integrated signal digital network (ISDN) circuit. It can also help facilitate the deployment of a line-shared loop. With respect to test access, lines can now be tested directly off the frame, fully automated or remotely. Unlike the narrowband test previously done through the classified switch, these tests can be performed on any pairs and are no longer restricted. Again, the testing can be full scale or full spectrum.

A new concept introduced to the industry is the switchable splitter illustrated in *Figures 8* and *9*. If a pair comes out to the customer premises, the voice would typically terminate in a classified switch and the DSLAM. Now those pairs can pass through a switchable splitter. The voice line goes to the loop just as it always has, but the splitter dynamically closes the loop and introduces an LPF. It also taps in to provide test access. All these activities take place remotely and can be potentially fully automated. The switchable splitter is important because it enables independent management of both POTS and ADSL services.

Figure 10 shows how the switchable splitter can potentially provide cross-connections if the DSLAM port fails. So rather than sending someone out to reconnect a pair to another DSLAM port, that pair can be dynamically switched to another active port on the DSLAM.

Providing port redundancy and also addressing service assurance are keys to reducing downtime. When providing an SLA and QoS to a customer, downtime becomes critical, particularly when a remote terminal or digital carrier is involved. An automated system like the one illustrated in *Figure 11* virtually eliminates downtime because switching occurs in less than a second.

FIGURE 8

Switchable Splitter

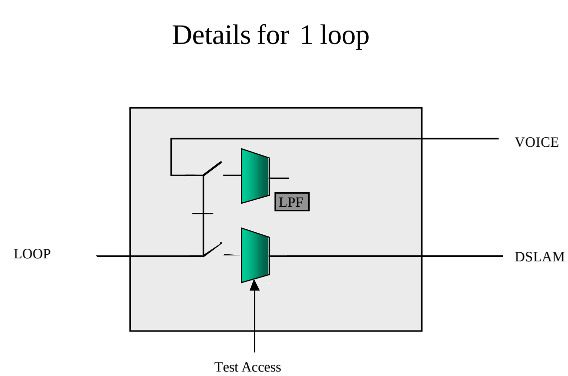

FIGURE 9

ADSL Service Activation

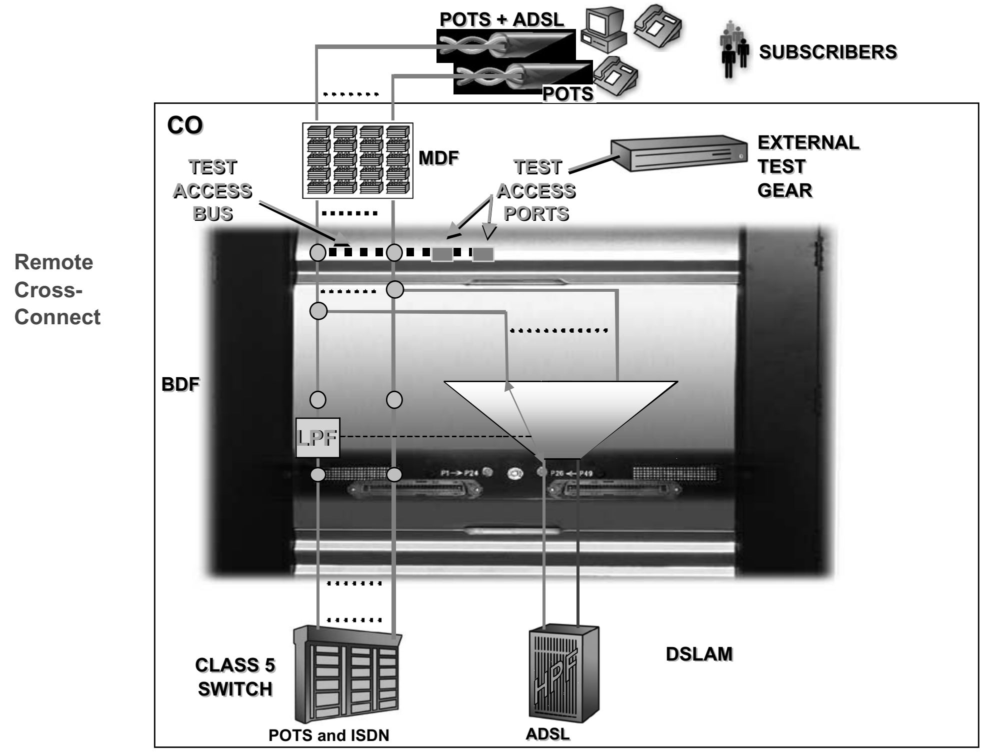

FIGURE 10

Protection/Fail-Safe Switching

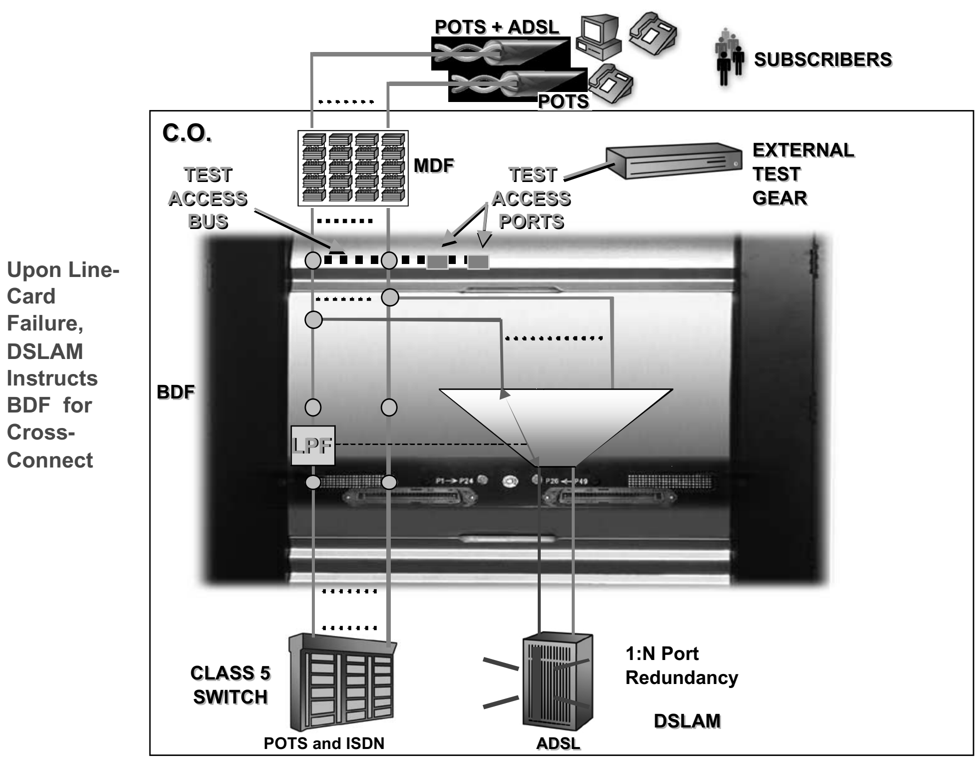

FIGURE 11

Protection/Fail-Safe Switching

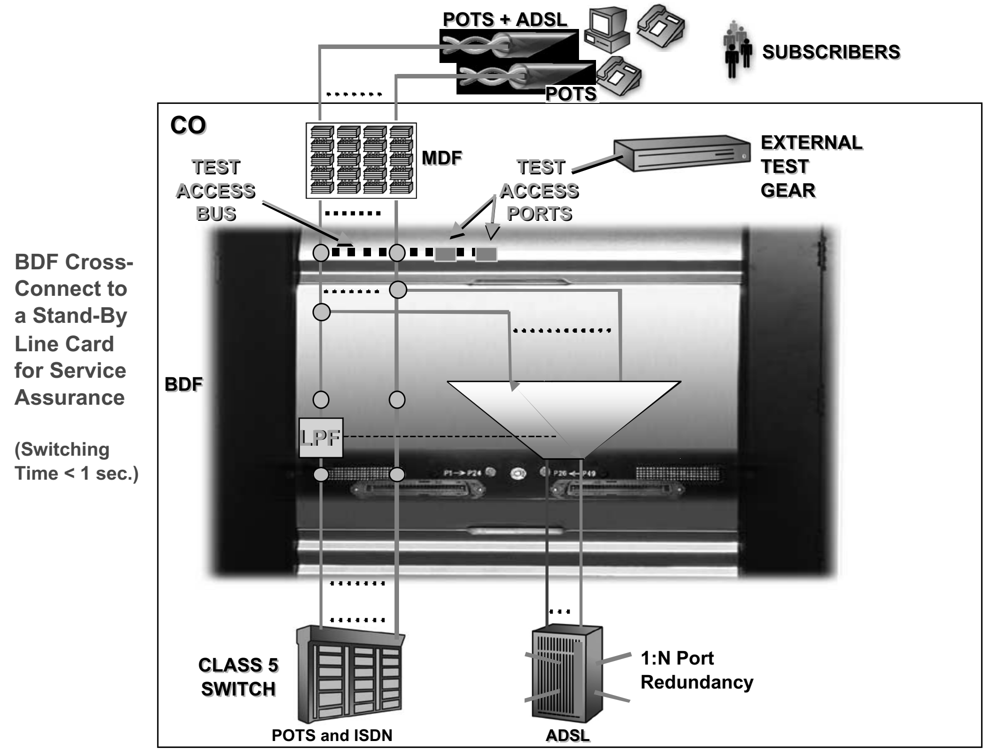

Figure 12 shows how some regulatory issues regarding delivering wholesale or other services can be addressed.

A carrier that introduces an LMS into its network can now facilitate and potentially automate the delivery of services to other carriers as well. With pre-wiring, it can now dynamically switch any pair to any potential customer. The carrier could be a CLEC or DLEC that is interested in deployment of services.

The NRIC unanimously recommended that all the carriers—CLEC or DLEC—no longer have to invest in their own test access infrastructures. The incumbent carriers would provide fair and reasonable test access for a fee. This is not policy yet; it is what the NRIC recommended to the FCC. However, It is the first time the carrier community has made a unanimous recommendation to the FCC. Pricing has not been determined and will perhaps be addressed at the state level.

The LMS discussed here can provide full broadband test qualification and limited or restricted access. When a customer orders a pair, a carrier can give the customer test access to those pairs through an LMS that has management or restricted test access.

Software and Remote Integration

Some software integration and some remote automation are needed to tackle some issues relating to good database records and the ability to clean up the database. *Figure 13* shows how integrated software and LMSs play a key role in doing just that.

Remote loop qualification reduces the time and cost necessary to build and maintain a qualified loop database that changes with the introduction of new services. Remote provisioning and test access facilitates faster service to new customers, thereby improving cash flow. It virtually eliminates the cost of manual intervention for service upgrades and churns, and reduces the trouble calls associated with initial provisioning. Traditionally, 15 to 35 percent of customers have made such calls—10 to 20 percent because of misconnections or no connections to the access network; 5 to 10 percent because of incorrect circuit length; and 0 to 5 percent because of undetected load coils.

Ongoing loop management reduces the cost of accuracy. Many carriers assume the costs of 5 to 10 percent of errors. With 10,000 lines per week, an error rate of 10 percent means 1,000 receive bad service or perhaps no service at all. Those are the types of issues that loop management can tackle.

BDF Benefits Summary

Some BDF benefits include the ability to do the following:

- Scale DSL deployment, connecting customers rapidly and remotely
- Meet specific requirements of a remote environment

Other benefits include the ability to do the following:

- Restore service rapidly during outages. Fast switching time allows on-the-fly maintenance and protection

FIGURE 12

Wholesale or Other Services

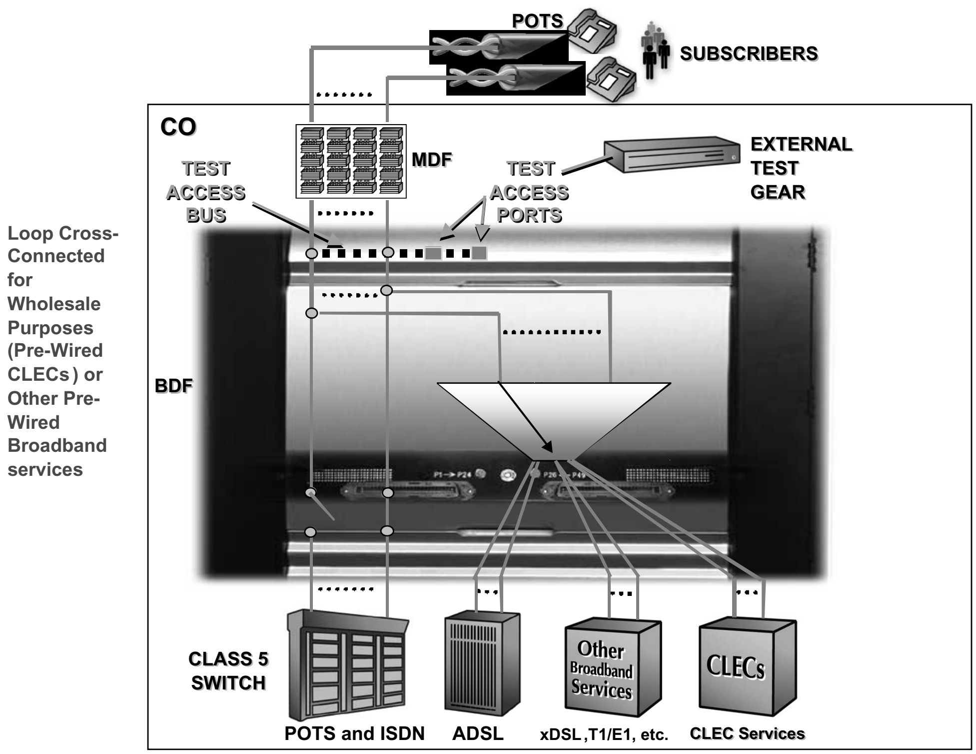

switching. During power loss, BDF maintains the circuit connections.
- Support multiple and redundant key components. BDF has no single point of failure.

FIGURE 13

Software Integration

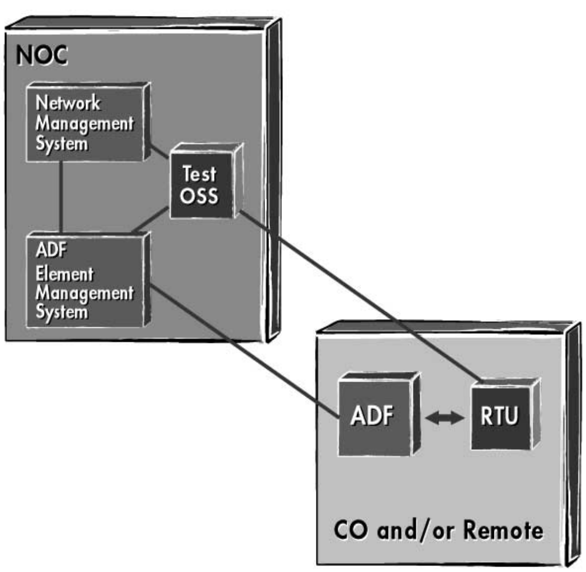

- Provide for service churn and upgrades through passive electromechanical processes transparent to protocols.
- Perform loop qualification through testing, monitoring, loop-back features, and integration to existing operations support systems (OSS).
- Simplify access management by keeping configuration records and easing the implementation of service changes. BDF's modular architecture lets service providers and customers both grow as they go.
- Provide a high-density solution that addresses cable management. The frame's compact switching gear is integrated within the connector block system and requires 50 percent fewer cables.
- Simplify local-network infrastructure deployment further with the remote version.

To these many benefits, add better, faster, and cheaper service with automated flow-through; lower operating and transportation costs; and higher density that enables integration of remote services. BDFs enable line sharing and VoSDL applications. They also reduce ILEC/DLEC dependence by resolving splitter-positioning issues and setting the demarcation point. By placing an adaptability feature in LMS, they reduce capital costs, and with protection switching embedded in LMS they enhance reliability and enable self-healing networks. By offering service on demand, they generate revenues.

Market Potential

The DSL industry is expanding rapidly, and the introduction of such capabilities as LMS contributes to that expansion. The expansion must continue in order to serve the millions of customers eager to use it.

TeleChoice estimates that 9.6 million U.S. businesses and residences will have DSL by 2003. Pioneer Consulting estimates 12.46 million. Revenues from residential markets will increase. Multiple additional applications will share the high-speed loop, providing integrated voice and data (VoDSL) at lower cost and, in many instances, replacing ISDN lines. Faster and better DSL deployment will increase revenue and address the cable threat. Some improvements described in this paper will help resolve problems related to class of service (CoS), QoS, and SLAs, and will increase revenues by as much as 60 to 80 percent per year. Compelling services are becoming available and will continue to roll out.

Learning More and Keeping Up

Many Web sites offer information on DSL

DSL Life (www.dsllife.com) is the DSL Forum's Web site for consumer education. It is a great resource that answers customers' questions about DSL.

TeleChoice publishes the latest announcements in the industry on its Web site, www.xdsl.com. On any given day, visitors might find 10 or 15 new DSL–related announcements.

The DSL Forum (www.dslforum.org) is the heartbeat of activity in the DSL industry. Most major service providers involved in DSL deployment today rely on and participate in it.

Notes

1. The author had this experience, even after working with the local telephone carrier for years to develop the technology and rollout some of the first DSL field trials.

"Battleground for the Consumer" or "Cable versus Telco"

Karl Rookstool
Vice President, Business Development
ADC Telecommunications

Introduction

The "war" for the residential customer is heating up again. The residential customer represents a very substantial bundle of telecommunications services ($150–$250) that both cable companies and incumbent local-exchange carriers (ILECs) are aggressively pursuing. This bundle of services is depicted in *Figure 1*.

Presently, the cable companies own the video market (90+%) while the ILECs own the voice market (90+%). The current battle is being fought over Internet access and the cable companies are winning. The next front will be cable's attack on voice. If Paul Revere were alive today and worked for an ILEC, he would be shouting: "The cable companies are coming! The cable companies are coming!" The big question is: What will be the telco's response?

Cable Threat

Revolution Round 1: Lasers and Digital Subscriber Lines
It was only 10 years ago that the first laser went into service in a cable network and initiated the hybrid fiber/coaxial revolution. The cable companies (and ILECs) immediately realized that new service opportunities were possible on this robust and scalable architecture. Beyond expanding channel choice, the first service on the cable companies' minds was voice. Around the same time, digital subscriber line (DSL) technology was emerging with the promise to dramatically increase capacity of twisted-pair copper wires. In these pre-Web years, the ILECs had video on their minds. With sabers rattling, many expected a war to break out between the multiple-system operators (MSOs) and the ILECs with video and telephony as the turf. But the first cable company thrust into voice and the initial trials of ILEC video were both essentially false starts and the expected battle never happened. The cable companies slow rolled telephony and the ILECs largely abandoned video. These early efforts were already on the wane when the next revolution round began.

Revolution Round 2: The World Wide Web and Satellites
The 1990s were only half over when two new revolutions showed up more or less at the same time: the Web and direct broadcast satellite (DBS) video. The Web opened a whole new application for DSL technology as well as an additional motivation for an upgraded hybrid fiber/coax (HFC) network. The cable companies began a five-year plant upgrade process that would expand channel capacity to compete with DBS and simultaneously enable two-way cable data. Meanwhile, the ILECs embarked on their own deployment plans for asymmetric DSL (ADSL) data service after abandoning video.

Current Competitive Status: Cable Ahead 2 to 1
Today, the cable companies and ILECs have delivered broadband data to more than 10 million North American subscribers. *Figure 2* shows some fairly recent quarterly subscriber additions for cable versus DSL. At least two cable-modem subscribers are turned up for every DSL subscriber. Why? There is no single explanation for this lead, but it is important is to understand the architecture and status of the cable-company HFC network in comparison to the ILEC twisted-pair network. A quick summary follows.

Superior Network: Capital Spent and Fiber Deep
The major cable companies are 80 percent to 90 percent complete with their upgrade effort. They have spent approximately $50 billion over five years to position lit fiber to within 3,000 to 6,000 feet of their residential subscribers. The HFC outside plant in place is highly scaleable going forward, and it is unlikely that large-scale upgrades will be necessary again anytime soon. As such, capital will increasingly shift to incremental, success-driven expenditures.

A common complaint about the HFC architecture is that upstream bandwidth is small relative to downstream. This is indeed true in a relative sense, but an operator has a number of relatively low-cost options to increase effective upstream bandwidth. Generally, they all follow a theme of reusing the available bandwidth through node decombining, radio frequency (RF) block conversion, etc. Importantly, these upgrade options can be carried out as needed and with *no additional fiber penetration.* Truck rolls, yes; trenching/boring crews, no.

Unified Standards and More Favorable Regulation
CableLabs drives cable-company technology consensus, writes end-to-end specifications, and certifies equipment.

FIGURE 1

War for Residential Customer

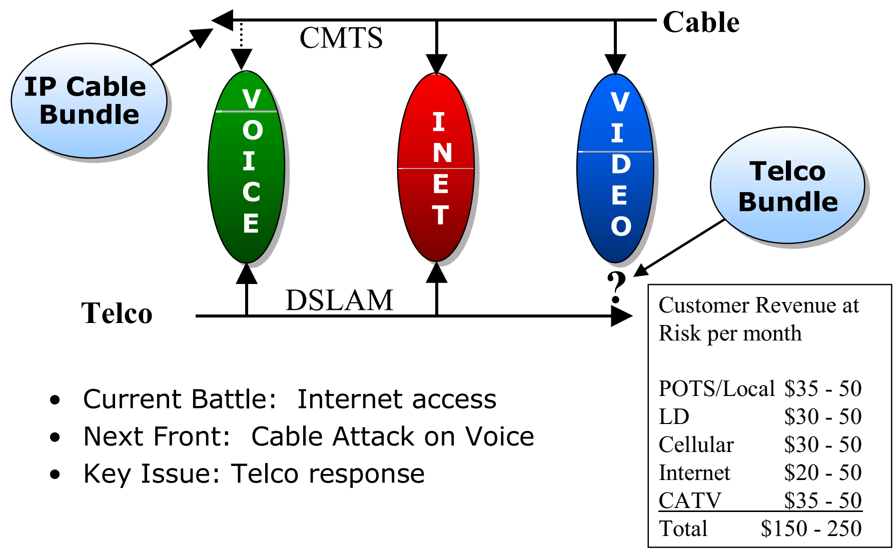

This unified front creates scale on many levels and drives down equipment and operating costs. It is an efficient and technically competent organization that has brought a great deal of order to the cable environment. After successfully launching best-effort data over cable service interface specifications (DOCSIS) 1.0, CableLabs is now certifying equipment for DOCSIS 1.1. DOCSIS 1.1 is the cornerstone around which the cable companies will introduce voice over Internet protocol (VoIP) and quality of service (QoS) into their shared access network. There is no counterpart to CableLabs for telcos. Telcordia Technologies, since its privatization, no longer provides this end-to-end interoperability industry function for the ILECs, although it may be contracted to assist on an individual company basis.

Cable companies are classified as Title VI service providers by the FCC, which affords them an advantage over ILECs that are classified as Title II incumbents. The Title VI designation does not require unbundling in the way that ILECs are forced to provide co-location facilities and physical access to loops. (AOL Time Warner is required to offer open access unbundling as a part of its merger terms and conditions. However, this is so-called packet unbundling as opposed to the more onerous co-location cages and physical unbundling.)

FIGURE 2

North America Subscriber Additions

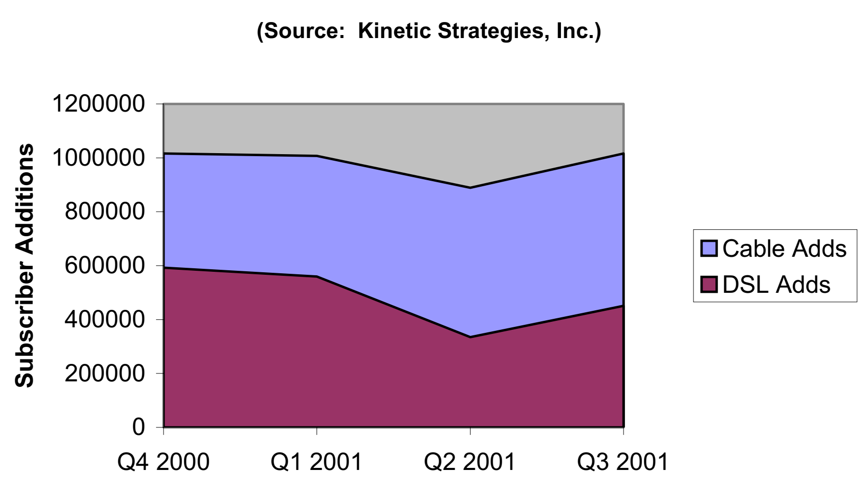

Telephony for Cable Companies Is Incremental and on Its Way

Even though the cable companies slowed their cable-telephony effort considerably during the last decade, they actually have more than 1.5 million voice subscribers today. From an architecture perspective, not much has changed: Telephony represents an incremental revenue stream on their existing installed plant. What has changed today is that DOCSIS 1.1 is just now being deployed, which will enable VoIP over the cable access network. The cable-industry effort to do telephony was not abandoned—it was simply delayed for a few years. Now, however, it will be built on the highly successful DOCSIS foundation rather than the proprietary circuit-switched systems of the past. Cable telephony never really went away, and now it is poised to come back with a vengeance. We expect cable television (CATV) early adopters to offer VoIP with DOCSIS 1.1 in 2H02. Mass deployment will commence in 2003. It is not as if the cable companies have the ILECs squarely in their gun sights. Rather, they need revenue growth and are targeting voice, video on demand (VOD), and other services to get growth. With second-line service clearly threatened and perhaps main line as well, what is a telco to do?

ILECs' Alternatives

The ILECs have several alternatives. They can continue business as usual and watch their second lines, then main lines, and ultimately entire service bundles migrate to the cable companies over time. They can bundle plain old telephone service (POTs)/local, long distance, cellular, and today's DSL Internet access with a targeted marketing approach to fight off the cable companies. They can attempt to level the playing field in the regulatory arena to enhance their bundling capability and burden the cable companies with some of their Title II requirements. While these last two alternatives will help the telcos to compete against the cable companies, they will still lack parity with cable's high-speed Internet access and video bundle. We believe the telcos must also modify their outside plant architecture to provide parity with the CATV broadband service bundle.

The telco's outside plant alternatives for fighting the MSO service bundle are as follows:

- Fiber-to-the-curb/fiber-to-the-home (FTTC/FTTH)
  - Passive optical network
  - Point to point (star)

- Utilize existing copper with DSL
  - Existing next-generation digital loop carrier (NGDLC), current distribution area (DA) architecture, current ADSL deployment plans
  - New fiber sweet spot, broadband digital loop carrier (DLC), and modified DA architecture

Not Fiber-to-the-Home

FTTH is the ultimate future-proof access technology. It has always been held up as the holy grail of ILEC access. Unfortunately, FTTH economics make it all but undeployable.

Not Enough Money

The basic reason is that the labor and material associated with a FTTH plant are not on any significant cost-curve decline. Construction costs have been steady even with advances in trenchless technology. Cable costs are already at commodity pricing. Even if dramatic reductions in optoelectronic equipment costs are possible in the future, the anticipated revenues from FTTH video/voice/data services do not justify the construction cost.

In a recent study, RHK stated flatly that "…it does not seem as if there will be any viable solutions for competitive residential fiber deployment in the near future." RHK studied a number of fiber architectures and found that all of them fell well short of viability for residential access. The most cost-effective FTTH architecture in their study was a 32:1 PON. To get a reasonable return on investment (ROI), the 32:1 PON architecture allows only $417 of equipment cost per home passed after PON construction. This equipment cost includes all network and subscriber termination equipment, routers, switches, servers, etc.—basically everything needed to deploy voice, video, and data but the PON itself. Note that the $417 target is exceeded today by the subscriber device *alone*, which ranges from $500 to $1,000 for an optical termination device. These numbers apply to new construction and not overbuild, where the economics would be considerably worse. Our analysis shows that a fiber/PON overbuild for a buried distribution area would cost about $2,000/subscriber.

For the sake of argument, let us assume that stupendous advances in technology and cost reduction enable the 32:1 PON architecture by driving per-home equipment costs below $417. Or alternatively, assume that some new service appears that dramatically changes the revenue side of the business model. In either case, a carrier would then have a reasonable business model and presumably could attract the capital necessary to deploy the plant and equipment. Even if the capital markets are willing to go along, there is another problem with FTTH that money can't solve: time.

Not Enough Time

It has been estimated that if all telco network construction was halted today and these resources re-applied to FTTH, a complete fiber buildout would take 15 years. It is a reasonably safe bet that the broadband winners and losers will be decided long before that time. There simply isn't enough time to deploy FTTH. To make matters worse, if a new service indeed appeared that made the FTTH business case attractive, then the MSOs would be in a better position to exploit it. They already have fiber to the node in their architecture and typically have lit optics to within 3 to 6 kft of a subscriber. This lead would save them considerable time and money in getting fiber to the home should service evolution support it.

Not Competitively Relevant

FTTH architectures are being trialed today in greenfield developments by major carriers and are also being deployed by some rural independents. Both of these deployments are, of course, highly subsidized. While FTTH trials are useful as a technology proving ground, there is simply no question that the FTTH architecture is not competitively relevant. There is not enough time or money for FTTH to make a difference in meeting the MSO threat. The simple facts are that 97 percent of the outside plant is already in the ground every year. At $2,000/subscriber, it is simply too expensive and would take too long to overbuild the existing plant with

FTTH. Telcos must find a way to optimize their most precious asset, which is the existing copper distribution plant.

Not Fiber to the Carrier Serving Area with Today's NGDLC and DAs

Project Pronto was SBC's plan to revamp their access architecture to offer nearly ubiquitous DSL service. Their stated goal was to offer a minimum of 1.5 megabits per second (Mbps) data service to more than 80 percent of their subscribers. Their architectural approach was to offer DSL out of remote terminals (RTs) and shorten loops to less than 12 kft. Although Project Pronto has been a part of SBC's DSL rollout to more than one million subscribers, they are in their own words, "fairly dramatically curtailing the build." Why?

Not Enough Deregulation, Bandwidth, or Revenue Opportunity

SBC blamed the regulatory environment as its primary reason for slowing down Project Pronto. The support for loop undbundling and co-location at RTs increased their costs considerably. SBC also claimed that the costs associated with active electronics in their neighborhood gateways were a large impediment to the business case. This is because the revenue potential and economic service life of the current platform and DA architecture is insufficient to provide an adequate ROI. The cost associated with remote electronics needs to drop by using next-generation technology, and they need to be able to evolve in a way that will generate more revenue per upgraded line by adding more services. The bandwidth target of 1.5 Mbps is too low for the scale and level of investment required in Pronto. This bandwidth is appropriate for Internet browsing, but this rate is likely to be more commodity oriented in the future.

Current NGDLC and Architecture Incapable of LAN/Video Speeds

The current generation of NGDLC platforms are incapable of local-area network (LAN)/video speeds. They have backplane bus-speed limitations, power distribution and consumption constraints, and heat dissipation issues. This means that the current planned expensive upgrades to existing NGDLC equipment only add Internet access as an application and will have limited economic service life. We estimate this service life to be two to three years. Today's generation of NGDLCs is fine for DSL services as long as a telco does not expect to be very successful and deploy many DSL services over the existing platform. In addition to the platform's limitations, there are also bandwidth limitations imposed by the current carrier serving area (CSA)/DA architecture. Loop lengths up to 12 kft will limit speeds to less than 1.5 Mbps. This combination of platform and architecture will never support LAN/video speeds. To truly compete with the MSOs on a full-service basis, the delivered bandwidth must be capable of video and that means an order of magnitude increase to 20–25 Mbps. Getting to these rates will require pushing fiber deeper into the network.

New Fiber Sweet Spot and Broadband DLC/DSLAM Required

We believe that a new fiber "sweet spot" and broadband DLC (BBDLC)/DSL access multiplexer (DSLAM) is required (see *Figure 3*). New platforms deployed must provide backplane speeds and software capable of supporting LAN/video speeds and services. The emerging generation of BBDLCs are designed from the get go to provide LAN/video speeds, QoS, and significantly improved densities. We also believe that DAs need to be reconfigured to new broadband DAs at distances less than or equal to 3 kft. This provides LAN/video speeds and competitive parity with the cable companies.

Capital Expenditure and Economics

On first blush, one would think that placing fiber within 3 kft and reconfiguring existing DAs to a new broadband design would significantly increase the capital investment and cost over today's plans of deploying DSL at the CSA (12

FIGURE 3

New BBDA

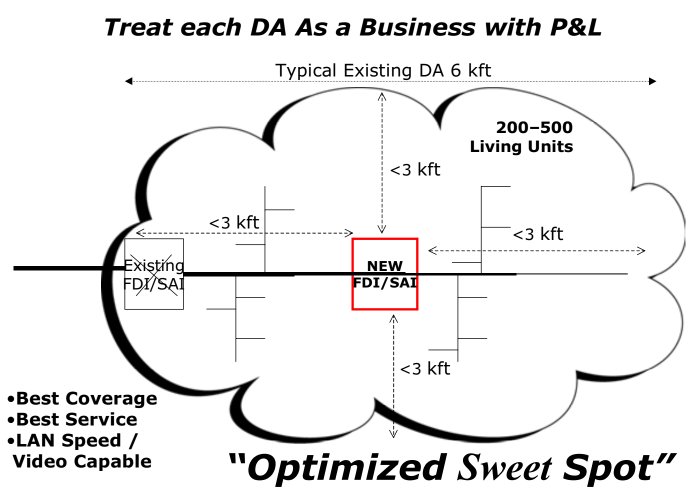

kft) (see *Figure 4*). This ignores the economic service issue of putting in an asset that will be short-lived versus one that is future-proofed. Based on a model of a CSA of three existing DAs and one new DA, totaling 2,000 POTS lines and 400 DSL customers, it would cost $120,000 in upfront capital with today's NGDLC and architecture. For the same model, using a BBDLC and reconfigured DAs to less than 3 kft, it would cost $220,000 in upfront capital.

However, the BBDLC and reconfigured DA would be LAN–speed/video-capable with an estimated service life of 10 years, while the NGDLC solution would only be capable of 1.5 Mbps Internet access with an estimated service life of three years. The higher $220,000 BBDLC/DA reconfigured capital investment has a lower monthly revenue requirement ($3,410) than the $120,000 NGDLC capital investment ($4,080). Not only is it capable of producing more revenue from LAN speed and video services, but it also costs $670 or 16 percent less per month. A comparison of the two alternatives is included in *Figure 5*.

If we assume that the average service bundle in this CSA is $150/month, the total revenue is $300,000/month (2000 subscribers x $150/month) and the annual revenue is $3,600,000 ($300,000 x 12 months). Surely protecting this kind of revenue with a more feature-rich, higher-speed capability is worth $100,000 more upfront capital, especially in light of the life-cycle cost reduction. The monthly revenue requirement per DSL customer would be $10.20 ($4,080/$400) with today's platform/DA versus $8.52 ($3,410/$400) for a BBDA. Thus, the cost reduction would be approximately $1.70/DSL customer/month.

Capital Expenditure Reality—"One DA @ a Time"

While having LAN–speed/video-capable DLC/DSLAM and DA architecture is the desired end game for telcos, reality is that they are unable to complete all DAs at once because of limited capital and resources. Therefore they must optimize their capital investment based on the following:

- Competitive threat
- Revenue opportunity
- Condition of outside plant
- ROI
- Feeder relief requirements
- Targeted marketing plans

In other words, telcos must plan "One DA @ a Time." The endgame for the ILEC is to optimize their copper distribution plant by having a BBDLC/DSLAM that is LAN–speed/video-capable within 3 kft of all customers. However, they cannot get there immediately, and " One DA @ a Time" is how to evolve to the endgame. The copper outside plant spans decades of time and technology. There are a wide variety of deployment rules, cable materials, and rehab scenarios used over the years. The intent here is not to cover these cases in exhaustive fashion but to give a set of general guidelines to be applied on a case-by-case basis. These recommendations are as follows:

- ILECs should only deploy DLC/DSLAMs that are LAN–speed/video-capable. Future-proof the network for greater revenue capabilities and economic service lives rather than upgrade aging NGDLC platforms.
- For all greenfield DAs and feeder relief projects, telcos should place BBDLCs within new BBDAs. These DAs should be designed with loop lengths less than 3 kft.
- Small, low-cost, line-powered, reusable DSLAMs should be used tactically to offer Internet access and retain existing customers when capital is scarce for low- and medium-priority DAs. These reusable DSLAMs should be moved and reused as BBDLC/DSLAM is brought to an area.
- Fiber should be placed to the new sweet spot and the DAs reconfigured for broadband when any construction occurs.
- ILECs need to be making their plans for value-added services over DSL, including voice, data, and video, as soon as possible.
- A targeted marketing program should be implemented to match the DSL service deployment.

FIGURE 4

2000 Living Units CSA Model

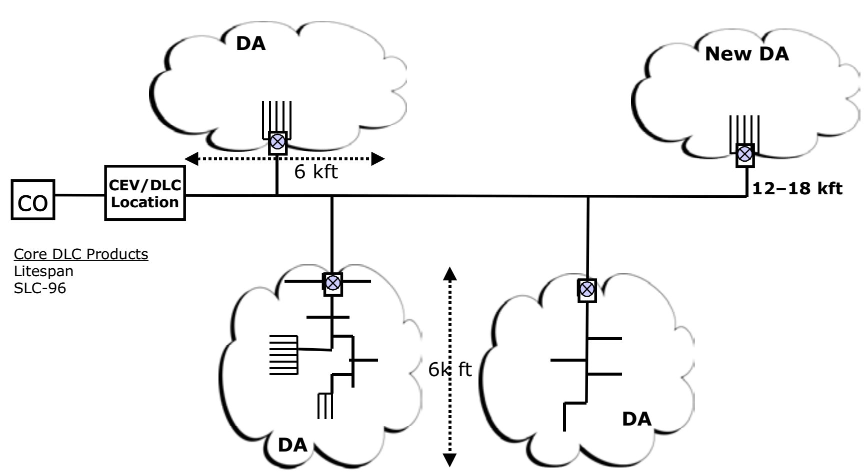

FIGURE 5

Capital Expenditure Comparison

2000 LU CSA, 4 DAs (3 Existing, 1 new), 20% DSL

	Today's Platform/DA	BBDA
Platform	NGDLC Upgrade	BBDLC
Loop Length	≤ 12,000′	≤ 3,000′
Speeds	≤ 1.5 Mbps	10–15 Mbps
First Cost	$120,000	$220,000
Service Life	3 years	10 years
Mothly Rev. Req.	$4,080	$3,410

- ILECs need to focus on the cable threat as their main regulatory story to achieve a level playing field. They need to be able to bundle their services to the same degree as the cable company.

Summary

The "war" for the residential customer is heating up and about to explode. The stakes are high with the entire bundle of telecommunications services, including local POTS, long distance, Internet access, cellular, and video ($150–$250) being at risk. For the ILECs, this time the cable threat is real. They aren't going away like the competitive local-exchange carriers (CLECs). The networks, technology, and ownership issues are converging, making the cable companies formidable competitors. ILECs have several choices. They can continue business as usual and watch the cable companies take market-share. They can aggressively bundle service capabilities and attempt to achieve regulatory parity. These actions will help to stave off the cable companies for a while but will not achieve parity with the cable companies' broadband services. We believe the outside plant architecture and DSL platforms must be modified to offer parity with the cable companies. FTTH is not the answer. It is both too expensive and too long to deploy. Today's DLC/DSLAM platform is not the answer. It will not have enough bandwidth or service capabilities to compete with the cable companies. It will result in expensive upgrades with very short service lives. The answer for the ILECs is optimizing their copper through "One DA @ a Time" planning. Fiber needs to be run to the new sweet spot, and DAs must be reconfigured for broadband. ILECs must be able to compete with cable companies on an equal bandwidth and service-capability basis. The cable companies are coming! The cable companies are coming!

A Software Library Perspective of DOCSIS 1.1: Enabling Cable Connectivity to Applications in a Diverse Environment

Mariano Schain
Chief Software Architect
Texas Instruments, Cable Broadband Communications (CBC)

Summary

The new 1.1 version of the data over cable service interface specification (DOCSIS) standard is targeted by the cable industry as the platform for consistent and reliable Internet over cable digital services through the use of sophisticated quality of service (QoS), security, and network management mechanisms. As a result, a variety of products are expected to embed DOCSIS 1.1 functionality, serving different applications such as home gateways, game consoles, telephony terminals, health and home monitoring systems, and more. The DOCSIS 1.1 component in such devices is viewed as a software library enabling the cable connectivity and control through a dedicated set of application program interfaces (APIs).

In this paper, the set of required software APIs is explored. A unified approach is taken, aimed at enabling the needs of the different applications with a generic collection of exposed functionality.

Introduction

DOCSIS 1.1 can be viewed as a framework—a set of tools that, if properly used, allows applications a reliable, secure, and efficient use of the cable medium. The functionality defined in the specifications is generic—no specific application is assumed; however, voice and video transmission needs are addressed.

With the aim to serve a variety of different applications, sometimes concurrently, a library approach as presented in this paper isolates the applications from the DOCSIS 1.1 details and hides them from the details of the native DOCSIS 1.1 messages and mechanisms. Instead, the DOCSIS 1.1 functionality and some common system resources are exposed through a dedicated set of APIs. This approach not only simplifies the system software architecture, but can also enable non-embedded applications to directly control the DOCSIS 1.1 library operation.

Packet cable is a set of specifications by CableLabs that builds on top of the DOCSIS 1.1 framework to implement an intellectual property (IP)–based packet telephony system. A natural example of a user application is the implementation of a packet cable embedded message transfer agent (MTA) on top of a DOCSIS 1.1 library.

In this paper, the DOCSIS 1.1 system and mechanisms are explained and the different set of APIs is presented. Throughout the discussion, packet cable is used as an example of a "real-world implementation" using the DOCSIS 1.1 API. This serves as an illustration of the way the APIs are used in a packet cable system to achieve efficient and reliable voice processing and signaling through a cable-modem system.

The DOCSIS Framework

Cable modem (CM) and the cable-modem termination system (CMTS) are the main entities in the DOCSIS network protocol. Several CMs, residing in the customer premises, are connected to a CMTS, which resides at the cable operators' head end through a hybrid fiber/coax (HFC) network. The HFC network is constructed in a tree-like structure, where the head end is the root and the CMs are the leaves.

The DOCSIS downstream (DS) channel, which carries information from the CMTS to the CMs, uses a typical television channel (6 MHz wide in the United States; 8 MHz wide in Europe) to carry Ethernet-like packets over a continuous digital Moving Pictures Experts Group (MPEG) stream. As the DS channel is shared, the channel bandwidth (about 40 Mbps in the United States) is distributed among all active CMs. Note that all the DS packets are received by all connected CMs. The Ethernet address is used by the CMs to filter out the packets that they need.

The DOCSIS upstream (US) channel carries Ethernet-like packets and uses the lower frequencies below the range allocated to television channels, which are prone to occasional interference. To cope with such situations, the US channel configuration, including bandwidth, rate, error correction, and other transmission parameters, is diverse and dynamically controlled by the CMTS. The transmission scheme is decided on a burst-by-burst basis. Again, the US channel is shared among several CMs, hence a multiple access mechanism is implemented and the channel bandwidth of up to 10 Mbps in the current specification is managed and allocated to active CMs by the CMTS. Note that US packets are received by the CMTS only. Packets from CM to CM always pass through the CMTS.

A DOCSIS domain may include several DS and US channels—paired accordingly to achieve the required network balance. CMs may be instructed by the CMTS to move from channel to channel as a load-balancing implementation or as a means to overcome channel-quality problems. A CM acts as a transparent bridge by forwarding packets that are received from the CMTS toward its local network interface (e.g., Ethernet, universal serial bus [USB], etc.) and vice versa. Packets that are destined to the CM, such as packets to the embedded simple network management protocol (SNMP) agent, dynamic host configuration protocol (DHCP) client, or other IP–based agents residing in the CM, are consumed by the CM and not forwarded.

As CMs reside in varying distances from the CMTS, the time division multiple access (TDMA) scheme implemented in the US requires subtle synchronization mechanisms. All CMs align to the CMTS clock, which is distributed through dedicated DS control messages. A ranging mechanism in which the CMTS instructs the CM on the time shift and power level to use in US transmissions is constantly active for every connected CM. This mechanism ensures that all CMs transmissions are aligned to a time base controlled by the CMTS and that all signals are received at the CMTS at approximately the same level, ensuring the ability of the CMTS to identify collisions.

The US channel is divided into time slots. A transmission interval is a group of continuous time slots. The CMTS allocates transmission intervals for different needs, including transmission requests, packet transmission, and ranging messages, and transmits the allocation to the CMs using a dedicated control message in the DS. Generally, transmission request intervals are not allocated to a specific CM. These allocated intervals are considered multicast and may result in a contention. Some of the intervals are unicast, such as a packet transmission interval. The US interval allocation messages, MAP messages, are transmitted in the DS by the CMTS in a timely manner. To illustrate the US transmission mechanism, consider a CM wishing to transmit a packet in the US. The CM will analyze the MAP messages in the DS until a multicast request interval is allocated by the CMTS. The CM will transmit its request, which contains the required transmission length, in the specified interval and will wait until a MAP message containing the requested allocation is received. Once the requested allocation or grant is received, the CM transmits the packet in the allocated interval. If more than one CM transmitted a request in the same interval, a back-off algorithm is implemented to solve the contention. To overcome the possible contention of subsequent requests, a CM may transmit a request embedded or piggybacked in a packet transmission.

While the best-effort scheme provided by DOCSIS 1.0 was sufficient for basic Internet access, it fell short to provide the needs of more sophisticated services that are unable to operate in the absence of guaranteed QoS. The QoS framework of DOCSIS 1.1 was targeted at exactly those types of services.

Four main service categories are supported by DOCSIS 1.1: unsolicited grant service (UGS), real-time polling service (rtPS), non–real-time polling service (nrtPS), and best effort (BE) service.

In a UGS flow, the CM is assured to receive from the CMTS fixed size grants at periodic intervals without the need to explicitly send requests, hence the service name. In addition to the grant size and the period, the tolerated grant jitter is also negotiated at service set-up. The main advantage in using a UGS is the reduced latency achieved by eliminating the need to go through the request-grant cycle for every packet. However, using a UGS is inefficient for applications that don't require a constant data rate over time. A flavor of UGS—unsolicited grant service with activity detection (UGS–AD)—is targeted at those exact applications (e.g., voice with silence detection). In a UGS–AD, once the CMTS detects flow inactivity through non-usage of grants by the CM, it starts sending unicast request opportunities at periodic time intervals. The CM can use the unicast request opportunities or polls to send requests once the flow resumes, avoiding the latency incurred by contention at multicast request intervals.

In a rtPS flow, the CM is assured to receive from the CMTS unicast request opportunities at periodic intervals. If the CM does not use the request opportunities, the CMTS allocates the reserved bandwidth to other flows, thus overcoming the inefficiency of UGS. In a nrtPS flow, the bandwidth is not guaranteed to the flow. The CM, however, is also allowed to use multicast request opportunities for the flow. The last QoS category, BE, defines the minimum traffic rate, which the CMTS must reserve, and the maximum allowed rate as the main service parameters.

Multiple data flows, each flow corresponding to a service and identified by a service ID (SID), may concurrently exist in a CM. A transmission request in the US and the corresponding grant includes the SID as the flow identifier. The CM and the CMTS negotiate the QoS for each flow upon allocation and dynamically as the service requirement changes in dedicated procedures. The QoS is then achieved by the implementation of sophisticated scheduling mechanisms in the CMTS. A classification function is applied to every packet. The flow in which a packet is transmitted is based on the content of the ethernet and IP header fields, allowing every application to receive a different service flow. The classification function may also indicate the suppression of the packet header payload header suppression (PHS), a mechanism that is useful for packets with semi-constant headers and short content, such as voice packets.

The DOCSIS framework includes some other mechanisms (e.g., privacy, secure software download, SNMPv3, etc.) that will not be discussed here and are not essential to the ideas presented in this paper.

In what follows, the set of APIs exposed by the DOCSIS 1.1 library is presented.

Dynamic Quality of Service (DQoS)

To support the establishment and configuration of the aforementioned QoS–related mechanisms, a set of media access control (MAC) management messages and procedures is defined. Service flows and the related classification parameters can be dynamically added, modified, or deleted through the use of such dedicated DQoS messages. However, as the DOCSIS 1.1 library is the only component that uses native DOCSIS 1.1 MAC messages, applications running on top of the DOCSIS 1.1 library, or users, are provided with a dedicated API that allows them to dynamically control the DOCSIS 1.1 QoS settings.

The DQoS API, through an abstraction layer, hides the DOCSIS 1.1 implementation details from the user. The challenge in defining the API is to balance the tendency to minimize the exposed complexity with the resulting limited control. For example, a decision has to be made regarding the way to expose the data structures that are internally used by the DOCSIS 1.1 library to represent the settings of a service flow (e.g., lists of classifiers, PHS rules, etc.). The API includes functions that enable the user application to initiate the establishment and configure the parameters of service flows, thus interacting with the CMTS through the CM, and also includes mechanisms for the user application to receive indications regarding asynchronous events (e.g., notification when a new service flow was established by the CMTS, indication of the successful acknowledgement by the CMTS to a service-flow addition request, etc.).

A packet cable application may use the DQoS API to establish UGS flows for phone calls. Indications regarding flow-establishment request status or requests initiated by the CMTS are also provided to the packet cable application through this API.

Performance Enhancements

According to the DOCSIS 1.1 specifications, every packet to be transmitted in the US should go through a classification process. The purpose of this process is to decide on the appropriate service flow, or the QoS parameters, that apply to the packet. However, sometimes the appropriate service flow is already known in advance, and bypassing the redundant classification process can result in better performance (i.e., less central processing unit [CPU] load, shorter delay, etc.).

Another example of a possibly redundant mechanism is the PHS expansion in the DS. In cases where the packet is to be immediately handed over to the embedded user application, it might not be required to expand the packet header, unless the user application needs some of the header information. In that case as well, bypassing the PHS expansion can result in improved performance.

The DOCSIS 1.1 library provides an API to support those types of performance enhancements, among others. In the US, the API includes a scheme to semantically identify the service flows that should be used in this special manner and a dedicated set of functions and supporting mechanisms inside the library that enable the user application to make use of those mechanisms. In the DS, the API allows the user application to indicate for every received packet whether it should be processed in the normal path or directly consumed by the user application.

By bypassing the US classification, a packet cable application can achieve lower delay on voice packet processing, hence at service-flow creation the required handles are set in order to enable such expedited handling by the DOCSIS 1.1 library. Also, PHS can be avoided for voice DS packets, as those packets are to be immediately consumed by the embedded voice processing entity, achieving an additional performance improvement.

Another improvement on the voice-processing delay is achieved by a packet-synchronization mechanism that is implemented in the DOCSIS 1.1 library and exposed by a dedicated API. The DOCSIS 1.1 library enables a user application or the packet cable application to get an indication on the actual timing of the US packet transmission grants. This allows the packet cable application to optimally schedule the voice sampling in order to achieve minimal voice delay.

Packet Processing

Some user applications implement a sort of packet filtering or manipulation (e.g., network address translation [NAT] or firewall). Those types of applications require that packets be handed over to them upon arrival through a network interface and prior to being sent through a network interface. To support such applications, a dedicated API is provided. The API enables the user applications to be notified with every incoming or outgoing packet. Note that there might be more than one application requiring notification. The user application may manipulate the packet content, indicate to the DOCSIS 1.1 library that the packet should be dropped, or allow further handling of the packet according to the DOCSIS 1.1 filtering specifications.

Shared Persistent Storage

User applications usually require a method to store persistent operational parameters. The DOCSIS 1.1 specification also requires that some parameters are persistently stored, and therefore a supporting mechanism is implemented as part of the DOCSIS 1.1 library using a flash-memory–based file system.

The persistent file system is made available to user applications through an API that allows them to allocate storage areas for private use. Storage-area content can be retrieved by the user application at start-up for initialization and may be updated upon a configuration change.

Logging

Local and remote logging mechanisms are defined as part of the DOCSIS 1.1 specifications. Those mechanisms allow the DOCSIS 1.1 library to report events to a logging server and to locally store recent events in persistent storage for later diagnosis.

Those mechanisms are made available to user applications through the library API. By utilizing this existing mechanism, the user applications can leverage the sophisticated logging options defined in DOCSIS 1.1 (e.g., error levels, throttling, etc.).

Management and Control

The operation of the user application may depend on the status of the DOCSIS connection and the DOCSIS operational parameters. To support such needs, an API to indicate changes in the cable link status is provided. This API allows any interested user application to register for such notification.

Another set of APIs allows the user application to retrieve the DHCP configuration as established by the DOCSIS 1.1 library as part of the DOCSIS 1.1 initialization mechanisms and to use the DOCSIS 1.1 trivial file transfer protocol (TFTP) client (e.g., for retrieval of user application configuration files).

A packet cable application uses the APIs to retrieve the MTA IP configuration parameters from the DOCSIS 1.1 DHCP client. It also uses indications regarding the cable interface link status, as this may affect ongoing phone calls.

Conclusion: The Benefits a DOCSIS 1.1 Library

Throughout the paper, the APIs of a DOCSIS 1.1 library were presented, with packet cable as an example of a user application. By encapsulating the DOCSIS 1.1 framework as a set of services provided by a library, user applications can utilize the DOCSIS 1.1 mechanisms and implement reliable and high-performance delivery of content over the cable infrastructure. The use of a library also enables efficient sharing of system resources (e.g., persistent storage and cable bandwidth) in case of concurrent user applications.

RPR Infrastructure for the Delivery of Deterministic Ethernet in the MAN

Bob Schiff
Senior Director, Marketing
Lantern Communications

This article examines resilient packet ring (RPR), an emerging technology and architecture for packet-based metropolitan-area networks (MAN).

Carriers' Services Portfolio—A Network Outlook

Service providers offer a range of bandwidth services to customers who need to connect multiple locations. A service provider's portfolio includes a mix of private-line, frame-relay, asynchronous transfer mode (ATM), wavelength and Ethernet services. The charts shown in *Figure 1* show the distribution of these services today and in four years. Looking ahead, we can expect growth in all of these services, but we believe the largest growth will be in the newest entrant, Ethernet. Scaling from kilobits to gigabits and being the most widely deployed data transport technology in the world, Ethernet has all the makings to become a core service offering for carriers worldwide.

Already a well-established fixture in corporate networks, Ethernet is quickly becoming the standard communications interface outside the enterprise as well. For example, hotels offer Ethernet modems in guest rooms supporting high-speed Internet access. At home, Ethernet is the standard interface between the home computer and the broadband (cable or DSL) modem. Home computers are now sold with Ethernet ports as part of their standard configuration. With so much momentum, it is only a matter of time before Ethernet becomes a strategic service offering for network operators.

Once we accept the premise that Ethernet services will be successful in the public network, it is then important to consider the network architecture that is best suited for its delivery.

The Importance of Ethernet

Why all the fuss about Ethernet? For one, it has an amazing lifecycle. When you consider the lifecycle of Ethernet to that of synchronous optical networks (SONET) and ATM, it dwarfs them both—and no one knows when the innovation will stop.

There are few technologies that can compare to Ethernet in its ability to evolve and adapt. Ethernet's longevity and durability are a testament to the elegance of the original concept and the effectiveness of the industry to evolve the technology to meet the needs of its newest users while preserving compatibility with its enormous installed base.

Ethernet has changed significantly from how it began over 20 years ago when it was introduced to the market as a shared, half-duplex 10 megabits per second (Mbps) local-area network (LAN) technology. It is now a full duplex technology that supports speeds from 10 Mbps to 1000 Mbps over copper and fiber media to distances of over 60 kilometers. The standardization of 10 Gigabit Ethernet is well underway, scaling this adaptable technology yet another order of magnitude.

Ethernet's flexibility and low cost, combined with its familiarity and ease of use, have enabled it to dominate the LAN landscape (over 95 percent). We expect this success to spill over into the wide area, and we expect that Ethernet-based services will be the successor to ATM and frame-relay services in MANs.

Figure 2 charts the innovation that has occurred over the past 25 years, starting with Bob Metcalf's patented invention back at Xerox PARC and moving forward to full duplex Ethernet, optical Ethernet, and Gigabit Ethernet. Standards that are expected in the next two years include 10 Gigabit Ethernet, first-mile technologies (for example, Ethernet over passive optical networks and Ethernet over DSL), and Ethernet aggregation and distribution solutions such as RPR.

Introducing an Ethernet access service is a natural choice for service providers. More than 95 percent of all data traffic originates and terminates as Ethernet. An Ethernet access service will seamlessly connect to a customer's access device that already supports Ethernet ports. This will greatly simplify the access solution by eliminating the cost and complexity of introducing another technology such as SONET, ATM or frame relay. Another factor in Ethernet's favor is its economies of scale. When thinking of Ethernet, you should not think in terms of hundreds of thousands of interfaces

FIGURE 1

The Changing Services Landscape

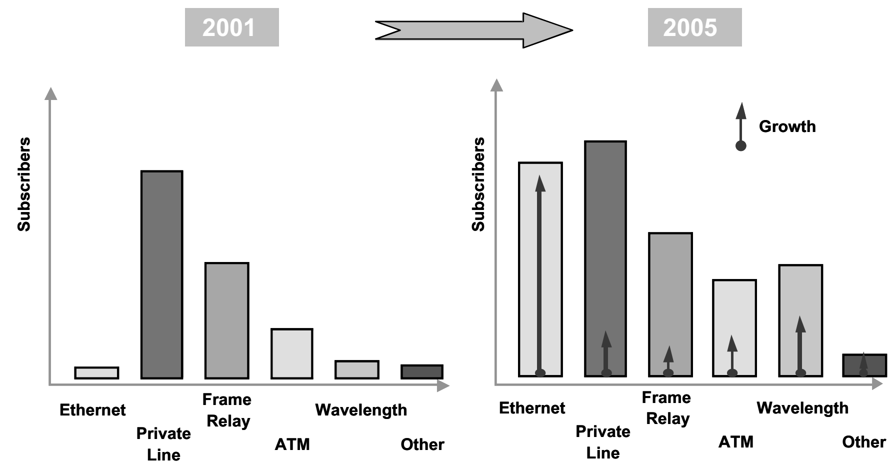

sold; you should think in terms of hundreds of millions of interfaces sold per year. These volumes translate to a large cost advantage over other networking technologies. Lastly, we must not overlook Ethernet's ubiquity and uniformity.

Ethernet is a universal technology that is defined by the same standards worldwide. Ethernet ports on equipment used in North America operate identically to Ethernet ports used on equipment used in Europe, Latin America, and Asia. This is a refreshing change from traditional telecom technologies that are defined by different standards in different parts of the world. Indeed, we believe that when life is found on other planets, they, too, will be using Ethernet!

Ethernet access is a scalable service offering that, when delivered over the right infrastructure, will permit the customer or the service provider to adjust the speed of the connection remotely. For example, an Ethernet access service delivered using a fast Ethernet port can be configured to support 20-Mbps throughput at one time and be quickly reconfigured to operate at 80-Mbps throughput at another time. This flexibility will allow customers to tune their access services on-demand to match their real-time traffic demands.

Ethernet will become the universal service jack for the 21st century. Once a service provider offers Ethernet access services with quality of service (QoS) capabilities, there are many revenue-generating possibilities. As shown in *Figure 3*, a provider will be able to offer integrated access solutions similar to the integrated access service provided today, but using all-packet techniques rather than time division multiplexing (TDM) schemes. Instead of using TDM timeslots to separate traffic and services, the differentiated Ethernet offering will use virtual LANs, Internet protocol (IP) addresses, and/or application port numbers to separate services and applications.

FIGURE 2

Continuous Innovation

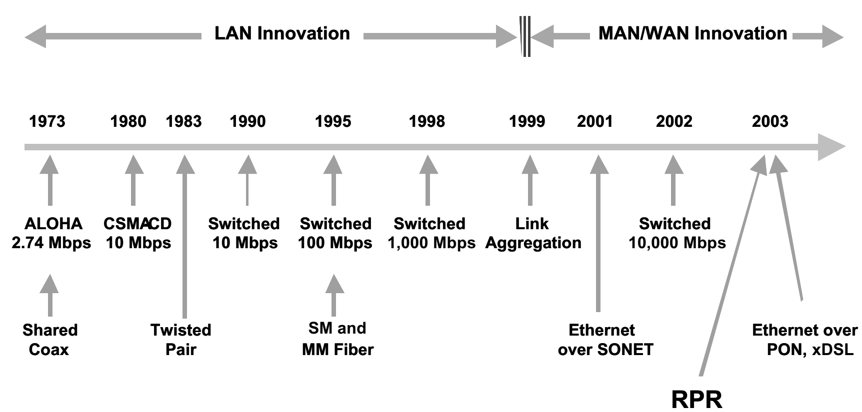

FIGURE 3

Enabling New Services

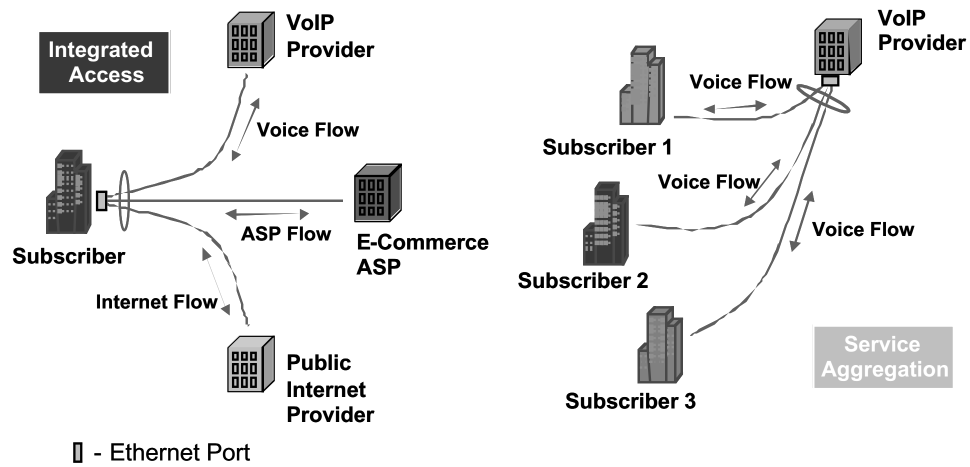

Service providers will use Ethernet as the interconnect technology of choice between networks, taking advantage of its capacity, low cost, and ubiquity for distributing and aggregating services to and from their end users.

Ethernet Service Infrastructure

But how best to deliver this Ethernet service? There is certainly no shortage of options—and this is the source of much confusion for customers and carriers. They hear from Ethernet-switch vendors, from packet-over-SONET vendors, Ethernet-over-SONET vendors, dense wavelength division multiplexing (DWDM) vendors, and now RPR vendors. It's very difficult for these customers and carriers to sort through all the marketing fluff and hype to identify what really makes sense in their given situation.

Figure 4 examines Ethernet delivery using SONET and DWDM. SONET's underlying TDM technology is not designed to accommodate the bursty and dynamic nature of packet applications. SONET's full-period circuit services have rigid boundaries that do not lend themselves to bandwidth sharing and dynamic bandwidth adjustment, both of which are important attributes of a flexible packet-switch network. Many of these next-generation SONET equipment add/drop devices use proprietary techniques for supporting or improving their data capabilities.

Metro DWDM systems that simply map Ethernet services to wavelengths are expensive and inefficient. Much like SONET, these metro DWDM systems lock up bandwidth and do not dynamically share unused capacity with other network users. This fixed allocation of network bandwidth results in stranded capacity and low utilization of network bandwidth. Indeed, there is much that can be done over a single wavelength of light before a service provider needs to consider introducing DWDM. DWDM systems often incorporate proprietary techniques for mapping data (e.g.,

FIGURE 4

Circuit-Based Solutions for Ethernet Service Delivery

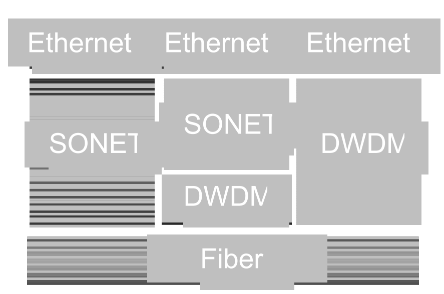

SONET/SDH

— Fixed payloads are not data friendly

— Burdened with limitations of TDM

— Inefficient for packet services

— Proprietary techniques for data support

Metro DWDM

— Not bandwidth efficient

— Expensive

— High operations overhead

— Proprietary

Ethernet) to their wavelengths. It is unlikely that the equipment of two DWDM vendors can talk to each other and pass Ethernet streams.

Figure 5 examines the Ethernet-switch solution set. For example, a provider may use Layer-3 Ethernet switches (i.e., routers). However, routers do not support an edge-to-edge Ethernet service and are complex and difficult to manage. Also, routers do not maintain QoS across multiple devices, a key requirement in a carrier-class metro infrastructure. Service restoration in a routed network can take tens of seconds—far slower than the 50-millisecond performance required of carrier-grade solutions.

Using Layer-2 Ethernet switches, the provider must still be concerned with the cost and complexity of managing multiple bridges. Ethernet switches, like routers, do not provide a mechanism for ensuring QoS across multiple hops. Service restoration in an Ethernet-switch network, again, can take tens of seconds—the convergence time for the spanning-tree protocol. The spanning-tree protocol actually breaks the ring, forcing a linear topology that does not take advantage of the ring for optimal path selection.

When we look at the different ways of delivering Ethernet service in the MAN shown in *Figure 6*, we see that there is only one architecture that incorporates a Layer-2 media access control (MAC) technology designed for the MAN. This architecture is RPR. RPR is an emerging technology being developed specifically for packet-ring environments to achieve maximum bandwidth utilization while preserving QoS commitments for packet applications such as Ethernet services. Neither SONET nor Ethernet switching addresses the need for a MAC layer designed for the metropolitan area environment (*Figure 4*). SONET employs Layer-1 techniques for bandwidth management and service protection. Ethernet switches rely on Ethernet bridging or IP routing for bandwidth management and service protection. Consequently, the network is either underutilized, in the case of SONET, or non-deterministic, in the case of Ethernet switches.

RPR addresses these shortcomings by defining a MAC layer optimized for metropolitan area networks that provides the resilience, deterministic service and manageability required of next-generation access networks.

Resilient Packet Ring Solution

RPR networks will be used as feeder rings in the MAN. These feeder rings, also known as collector rings, will be used in much the same way as SONET rings today, but for

FIGURE 5

Ethernet Switch Solutions for Ethernet Service Delivery

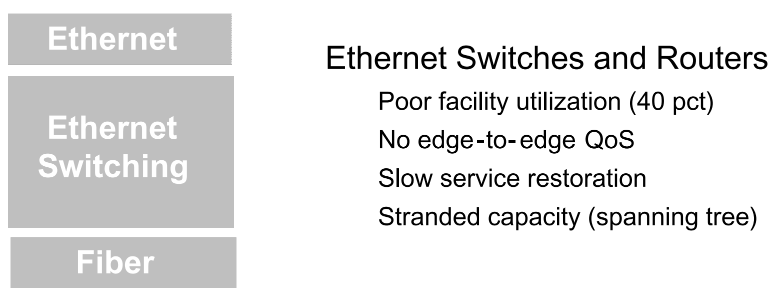

FIGURE 6

Ethernet in the MAN

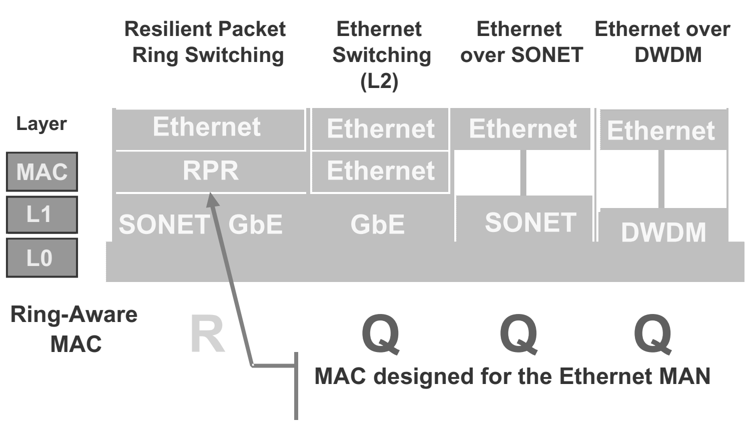

packet traffic. Packet access networks such as cable-modem systems and fixed wireless networks will feed RPR networks as shown in *Figure 7*. The RPR network aggregates the packet traffic from around the ring onto one or more highly concentrated links connecting to the network edge. At the network's edge, these links are the handoffs to the DWDM and router networks that switch the traffic to its final destination.

RPR standards are currently being worked on within the Institute of Electrical and Electronics Engineers (IEEE) 802.17 RPR Working Group. The time line for the first draft to be available is January 2002, with an adopted standard in March 2003. The RPR Alliance is a marketing alliance created by a number of RPR companies to promote RPR applications and to accelerate interoperability testing and deployment of RPR solutions.

Several key features distinguish RPR solutions from SONET and Ethernet switch solutions. RPR provides a ring-aware solution for MAN packet networks. It supports a bandwidth-management mechanism that ensures bandwidth commitments at a node or subscriber level. RPR incorporates a service-restoration method that guarantees service protection within 50 milliseconds of a fiber cut. RPR also defines a closed-loop congestion-control technique that enables maximum utilization of the ring capacity while preserving service-level commitments.

The end result is a packet-optimized solution that fully utilizes fiber capacity, provides SONET-class resilience, and supports thousands of packet applications, including metro Ethernet services, each backed by service-level commitments. To illustrate this advantage, *Figure 8* contrasts the number of metro Ethernet services, each configured with a

FIGURE 7

Using RPR for Traffic Distribution and Aggregation in the MAN

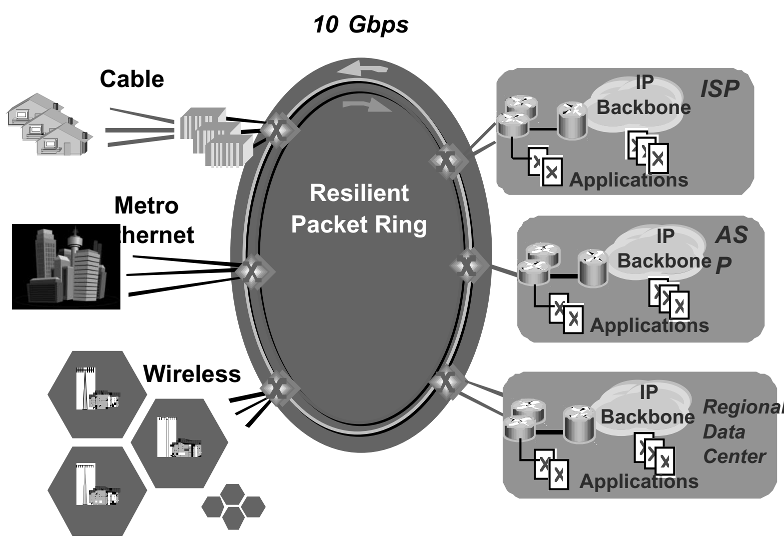

FIGURE 8

RPR's Disruptive Advantage: Ten times the service at half the cost

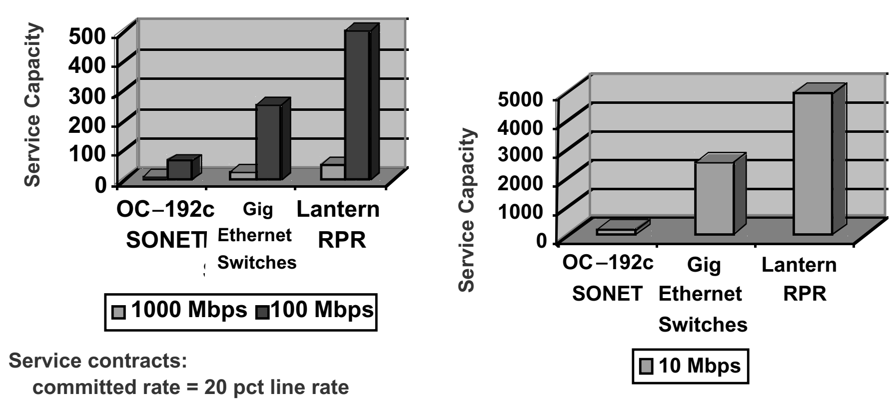

committed bandwidth equal to 20 percent line rate (e.g., 20 Mbps for an Ethernet service delivered using a 100 Mbps Ethernet port), that can be provisioned on a SONET OC-192c 2-fiber ring, a dual 10 gigabit Ethernet ring (Ethernet switches), and a dual 10-Gbps RPR.

The RPR solution supports over double the number of Ethernet services than Gigabit Ethernet switches, and over 10 times the number of Ethernet services than OC-192c SONET systems.

Perhaps most important is the solution cost. A network model consisting of five add-drop locations (four tributary sites and one hub site) was priced out for SONET and RPR. The service contract used in the example is a Gigabit Ethernet-based offering with a 500-Mbps committed bandwidth—i.e., a 500-Mbps committed rate and a 1,000-Mbps peak rate (see *Figure 9*). The SONET on the left prices out at approximately $900,000 for a dual OC-192c ring with four add/drop multiplexers (ADM). The RPR solution—with the same number of service contracts, but only utilizing 40 percent of the ring capacity—prices out at approximately $500,000, or nearly half the cost of the SONET.

Conclusion

In this paper, we have contrasted several approaches to delivering Ethernet services in the MAN. Circuit-based solutions, characterized by their static-full period, TDM-based connections, are neither cost-effective nor bandwidth efficient for Ethernet service delivery. Ethernet-switch solutions are not designed for the carrier environment and do not provide the edge-to-edge QoS or rapid restoration required of metro infrastructure.

RPR is a technology and architecture that is being developed for packet-based MANs. Leveraging the efficiencies and economics of packet technologies and designed to deliver SONET-class performance and reliability, RPR is the most scalable and cost-effective architecture on which to deploy Ethernet and other packet services on metropolitan fiber rings. RPR networks will play an important role in the evolution of the public network, providing a reliable, scalable, and efficient packet edge for the emerging optical Internet.

FIGURE 9

Economic Advantage of RPR

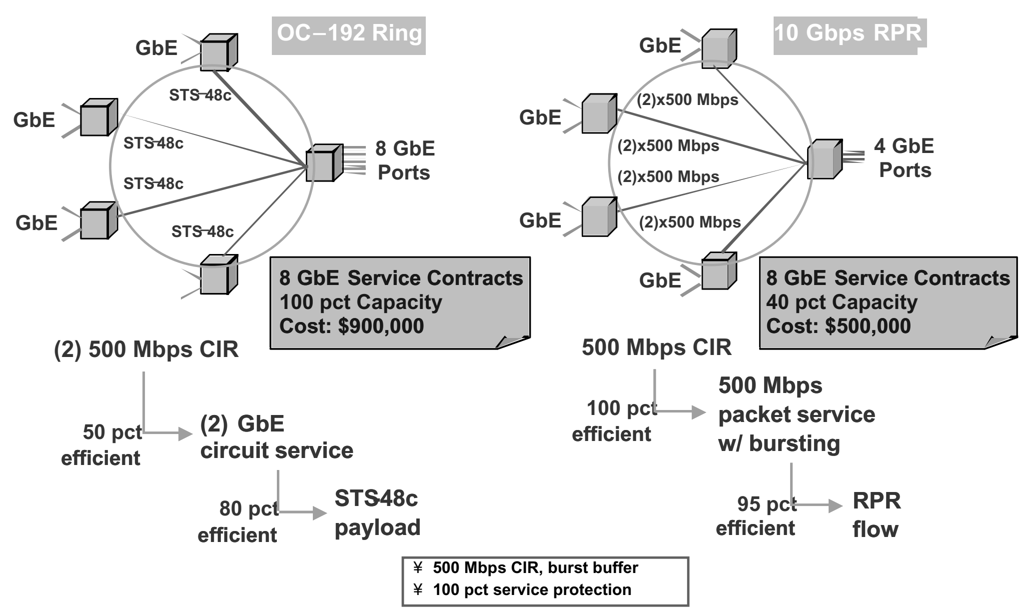

Fibre-Channel Extension Services

Paul Schoenau
Enterprise Product Marketing Manager
Akara

Today's optical-services market is characterized by what has become known as the "optical chasm"—the barrier that is separating enterprise need for services and the ability of service providers and carriers to deliver cost-effective, business-critical offerings. This paper explains how emerging technologies can bridge this optical chasm, finally delivering cost-effective storage services while reducing capital and operations costs.

As today's enterprise marketplace continues to demand 24/7 data-center operations, information technology (IT) managers are under increasing pressure to support the networking of more and more high-availability applications such as disk mirroring, remote tape backup, and server clustering across the optical metropolitan-area network (MAN) and wide area network (WAN). To support these applications, many enterprises have created storage-area network (SAN) architectures built from various combinations of servers, disk arrays, tape silos, switches, and networking equipment, typically connected with Fibre Channel links.

Current solutions to network Fibre Channel have relied on link extenders (LE) and dense wavelength division multiplexing (DWDM) (see *Figure 1*)—transparent solutions that meet only the basic connectivity requirements of carrying the 1s and 0s of data-center protocols.

Connectivity alone, however, does not meet the networking requirements of today's data centers. To deliver a truly efficient, high-performance data-center operations solution, Enterprise chief information officers and network managers are increasingly demanding full, turnkey services with capabilities that go beyond today's solutions and provide seamless integration of the SAN and MAN networking environments. These capabilities include the following:

- Support of data-center applications over distance, including support of Fibre Channel and the problems that distance introduces;
- Detailed optical-service visibility, control, and security; and
- Scaling of optical networks to accommodate growth.

To address the emerging needs of data-center networking, service providers and carriers are turning to protocol-aware optical platforms that surround the optical network, providing reduced costs and opening the door to diverse new service offerings. The use of these types of protocol-aware platforms can equip the optical layer for the delivery of storage networking traffic by providing the following benefits:

- Optimized use of networking resources:
  - Up to 90 percent more efficient use of the network;
  - Bandwidth delivered that matches the application's requirement; and
  - Zero packet loss.
- Enhanced optical services networking:
  - Improved service visibility;
  - Service security; and
  - End-to-end service management.
- Interoperability with existing equipment:
  - Non-intrusive Fibre-Channel interoperability;
  - Improved Fibre-Channel transport performance; and
  - Easy integration into existing synchronous optical network (SONET) and DWDM infrastructures.

Also, By using native-protocol flow control to combine packet-based flows of Fibre Channel and Gigabit Ethernet into a standard SONET payload that can be carried on dark fibre, SONET, or a DWDM optical network, these platforms can bridge the data center and the optical MAN and WAN, providing:

- Optimized control of the optical network according to individual application requirements, such as latency and throughput;
- Visibility of performance metrics and usage statistics; and
- Efficient, overhead-free data multiplexing.

Time division multiplexing (TDM) technology can deliver deterministic performance and security isolation demanded by enterprises' storage applications. A clear demarcation point between the enterprise and the network will result in greater service reliability, network accountability, and lower operational cost. The result is services beyond simple optical connectivity.

Key Benefits

Optimized Use of Networking Resources

Up to 90 Percent More Efficient Use of the Network

A protocol-aware platform can eliminate the need to forward unnecessary protocol "white space" which, in a typical application, could consume as much as 90 percent of the

FIGURE 1

Current Solutions to Network Fiber Channel

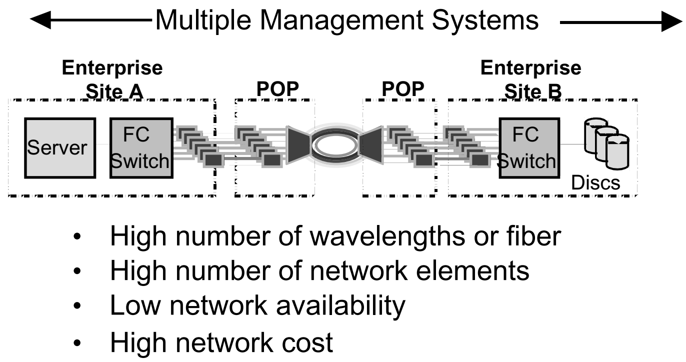

available bandwidth (see *Figure 2*). By replacing white space with control and data packets from multiple services, more "real data" is carried per unit of fiber deployed, and bandwidth use is optimized.

Bandwidth Delivered to Meet Application Requirements
When bandwidth is provisioned to support the actual application requirements, the need to over-provision the network is eliminated. This process simplifies network engineering, provisioning, and management, and unused bandwidth can then be allocated for additional revenue-generating services.

Zero Packet Loss
The allocation of a specific amount of bandwidth end-to-end specifies a deterministic path through the network per service, without any loss of data.

Enhanced Optical Service Networking
Improved Service Visibility
End-to-end visibility of Fibre-Channel optical data-center services can also provide detailed end-to-end service information including the following:

- Performance metrics;
- Latency measurements;
- Flow control monitoring; and
- Link statistics.

Protocol error checking at each stage in the service circuit enables rapid fault isolation. Enhanced service-level agreements (SLA), with detailed statistics that can be measured and reported, can also be created (see *Figure 3*).

Deterministic Security
Data-center traffic needs for guaranteed security and high performance require the network to provide end-to-end service isolation with no sharing or contention for bandwidth resources.

End-to-End Service Management
By collecting, processing, and correlating detailed service information, and by consolidating that information into a single, end-to-end service view, interoperability with other network operational support systems can be achieved. The ability to manage and control high-speed storage services can also result from this form of network management.

FIGURE 2

White Space

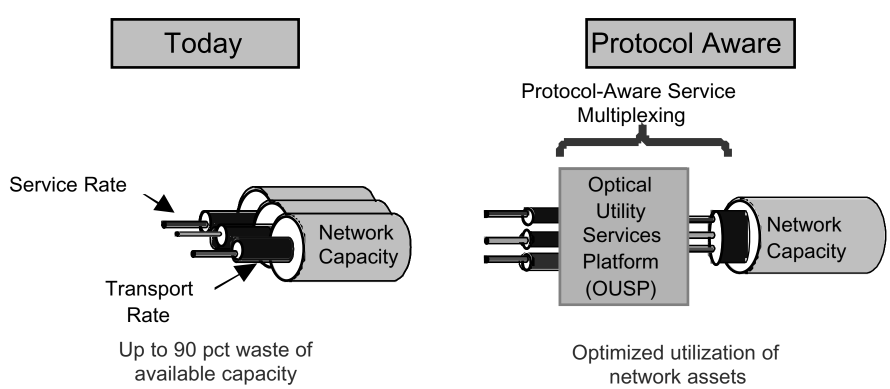

FIGURE 3

Enhanced SLAs

Today

Basic SLA

Mean Time to Respond	30 min
Mean Time to Repair	4 hours
Service Availability	99.80 pct
Network Monitoring	24 x 7

Protocol Aware

Basic SLA

Mean Time to Respond	30 min
Mean Time to Repair	4 hours
Service Availability	99.80 pct
Network Monitoring	24 x 7

Enhanced SLA

	Gold	Silver	Bronze
Client PMs			
CRC Errors	☑	☐	☐
Latency Monitoring			
Roundtrip	☑	☐	☐
One Way	☑	☐	☐
Average	☐	☑	☐
Real-Time Statistics			
Frames Transmitted/Received	☑	☐	☐
Bytes Trans/Received	☑	☐	☐
Pct Utilization	☑	☐	☐
Protection Options			
Pct Guaranteed	☐	☑	☐
Pct Unprotected	☐	☑	☐
Pct Pre-emptible	☐	☑	☐

Leveraging Existing Assets

Non-Intrusive Fibre Channel Interoperability

By interoperating with the data-center protocols at Layer 2, the values of efficient multiplexing, protocol error checking, and native flow control can be leveraged without impacting performance and operation of the link.

Improved Fibre-Channel Transport Performance

A flow-control system can be set to smooth and pack frames into the provisioned line bandwidth, maintaining the utilization at nearly 100 percent. For latency-sensitive applications, flow control can be set to "wire speed" latency, making the packet latency equivalent to a piece of wire with the tradeoff of less efficiency on the line utilization.

Easy Integration into Existing Optical Infrastructure

By multiplexing packet-based flows into a standard SONET, the payload can be designed to interoperate with existing infrastructure, whether SONET, DWDM, or dark fiber.

Summary

Together, these benefits—optimized use of networking resources, enhanced optical-services networking, and interoperability with existing equipment—will enable today's enterprise data centers to quickly and cost effectively meet rising demand for high-availability storage networking applications, and implement solutions that maximize the performance, security, and efficiency of carrying Fibre-Channel links across the MAN and WAN.

The Carrier-Class Gateway in the Voice over Broadband Access Evolution

Mike Schroeder
Vice President, Technical Services
General Bandwidth

Issues with Initial Deployment of Voice over Broadband Access

Initial deployment of voice over broadband access focuses on broadband loop emulation services (BLES), a joint standard from the ATM Forum and the DSL Forum. BLES is based on asynchronous transfer mode (ATM) virtual circuits over digital subscriber line (DSL) between a customer premises integrated access device (IAD) and a voice gateway. The gateway, in turn, interfaces to a Class-5 switch.

BLES provides no routing intelligence or call features, all of which come from the Class-5 switch. BLES, therefore, facilitates early deployment of voice over DSL (VoDSL) by reducing the functionality required in the IAD, thus reducing the unit cost of the IAD—an important early-deployment feature, considering an IAD is required for each customer premises.

The following sections describe issues that can arise in the evolution of a BLES architecture to a distributed voice over Internet protocol (VoIP) architecture.

Limited Intelligence in IADs

Over time, the deployment of VoDSL using low cost, limited intelligence IADs results in an installed base of IADs with limited processing power, limited memory, and no VoIP call signaling capabilities for providing softswitch-derived plain old telephone service (POTS).

Although these IADs successfully provide BLES capability, they cannot function in a VoIP–based softswitch architecture and must be upgraded or replaced—a very expensive proposition in a large-scale BLES deployment.

Customer Use of VoIP

BLES uses voice over ATM (VoATM). However, voice can enter the protocol stack at three different points: VoIP, VoATM, and channelized VoDSL (CVoDSL) (see *Figure 1*). CVoDSL does not take advantage of bandwidth savings resulting from silence removal since it uses fixed bandwidth DSL channels. Additionally, it requires the carrier to continue to expand the time division multiplexing (TDM)–based Class-5 network and is therefore less likely to be used. Internet protocol (IP) is the dominant technology in the customer premises and is the technology that carriers view as a better fit for their long-term strategy for service delivery.

Deployments of IP phones, IP private branch exchanges (PBXs), and VoIP functionality in desktop computers, laptop computers, and personal digital assistants (PDAs) require signaling capabilities such as session initiation protocol (SIP), media gateway control protocol (MGCP), and media gateway control (MEGACO)/H.248 that are not available in Class-5 switches.

The incumbent local-exchange carrier (ILEC) network must provide access to call-control servers (softswitches) that implement these protocols in order to avoid losing revenue from customer-premises implementations.

This is the old Centrex/PBX issue revisited. VoIP allows great flexibility in its distributed architecture for call control, allowing either the customer or the carrier to provide the necessary functionality. Carriers seek to provide as much functionality and network intelligence as possible in order to offer a wide range of revenue-generating services and to prevent their networks from being used for low-margin bit pipes, leased lines, and virtual private networks (VPNs).

It is important that carrier-based gateways can transport VoIP traffic from both the IP PBX and VoIP Centrex scenario.

Routing of Voice and Data Traffic

When customer traffic arrives over the DSL access network, an ATM switch aggregates it. The voice traffic is forwarded to the voice gateway, and the data traffic is forwarded to one of several networks.

FIGURE 1

Alternatives for DSL Transport of Voice

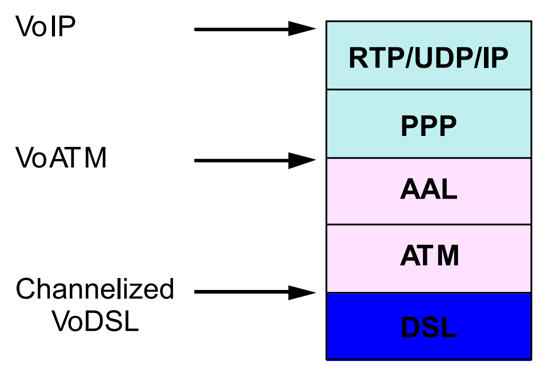

- An ATM switch for transport to an Internet service provider (ISP) or for access to the carrier's public ATM service
- A router providing access to the public internet or a managed IP network
- A frame-relay switch for transport to an ISP
- A label-switched router (LSR) for transport over a multiprotocol label switching (MPLS) backbone

This results in multiple pieces of data switching equipment in the central officer (CO), resulting in issues of the following:

- Equipment footprint, power, and physical space in the CO
- Interoperability among equipment possibly coming from different vendors
- Management of multivendor equipment resulting in several different element-management platforms
- Increased cost associated with purchase of many different pieces of equipment
- Routing delay resulting in end-to-end quality of service (QoS) issues and poor voice quality

Interworking between Class-4 and Class-5 Softswitches and Gateways

Class-5 switches implement Class feature software, whereas Class-4 switches do not need that capability. Class-5 switches need a range of interfaces to the customer, including POTS, integrated services digital network (ISDN), xDSL, T1/E1, T3/E3, STS–3/STM–1, whereas Class-4 switches generally just need a range of STS–n/STM–n interfaces. As the software-intensive functionality and the diversity of physical interfaces of a Class-4 switch is less complex than that of a Class-5 switch, it is the Class-4 switches that are being replaced by softswitches and gateways before Class-5 replacement. Moreover, the vendors specializing in Class-4 replacement are different from the vendors specializing in Class-5 replacement.

During an evolution to an end-to-end packet voice architecture, it is imperative to ensure interworking between Class-4 and Class-5 softswitches and gateways.

Evolution from MGCP to MEGACO

Initial deployment of softswitch architecture in the public network is using MGCP, primarily due to the simplicity of its implementation. This raises issues of functionality that is not available in MGCP or is provided more flexibly by implementing the MEGACO standard.

Interworking with Customer-Premises H.323

At the customer premises, the H.323 packet voice technology has the largest installed base. In the public network, MGCP currently is the most prevalent packet voice technology. As a result, it is imperative that H.323 and MGCP/MEGACO interwork in public network softswitches.

Provision of Emergency Services

There are two issues related to provisioning of emergency E-911 service: customer-premises equipment (CPE) and network equipment.

Customer-Premises Equipment

Current POTS lines provide power to customer premises analog phones, and PBX customers generally incorporate UPS backup to ensure that the subscriber can make emergency calls in the event of a loss of electrical power.

Packet voice at the customer premises and in the public network must also ensure that the customer can continue to make emergency E-911 calls in the event of a power loss at the customer premises.

Network Equipment

Network equipment must be able to provide a number of features specific to E-911 calls, including the following:

- Emergency ringback if the call is disconnected
- Calling party on/off hook status provision to E-911 operator
- Route diversity
- Multifrequency tone provision as required by tandem switch from which emergency service is provided

Provision of Wiretap Capabilities: CALEA

The 1994 U.S. Communications Assistance for Law Enforcement Act (CALEA) requires carriers to provide wiretap capabilities in cases authorized by court order, and reimburses them for the associated cost. CALEA implementation for packet voice is described by the Federal Communications Commission (FCC) in 97-213.

Wiretaps authorized by law enforcement authorities are implemented in Class-5 switches for analog POTS lines and TDM access facilities. Packet voice must provide equivalent functionality.

Conclusion

There are still significant issues related to the migration of voice services from a Class-5/BLES architecture to a fully packetized distributed softswitch architecture.

The ideal gateway platform is designed to allow carriers a graceful migration path from a Class-5 derived service, such as BLES, to a fully packet-based softswitch by supporting both services concurrently. Additionally, as new access architectures are deployed, such as ATM passive optical network (APON), Ethernet passive optical network (EPON), and LRE, the G6 platform can continue to provide both Class-5 and softswitch-derived services. The key features that allow the platform to provide this migration are as follows:

- Concurrent support of VoIP and VoATM
- Support of MGCP and MEGACO standards for softswitch deployments
- Support of softswitch control of BLES–based IADs
- IP routing and MPLS support for IP–based deployments
- QoS interworking between IP, ATM, and MPLS
- Scalable ATM/IP switching architecture for access aggregation

A graceful evolution is a complex undertaking, and the carrier-class gateway is at the center of that evolution.

Deploying All-Optical Access Networks

Shane Shaneman

Market Manager, Access Solutions and Business Development

Corning Cable Systems

This article examines the emergence of the passive optical network (PON) cloud and the dynamics and drivers behind the deployment of fiber further in the access network. It will examine the primary PON architectures deployed to date, discuss the advantages and disadvantages that each represents to a carrier's network strategy and business case, and concludes with a detailed cost analysis of passive-infrastructure cost drivers for PON deployments.

Emergence of the PON Cloud

Over the past 12 to 18 months, much attention has been given to the development and deployment of all-optical access networks to meet the continually growing demand for bandwidth in the local loop. Many active equipment vendors have developed equipment that is focused on enabling the delivery of broadband services and applications such as fiber-to-the-home, fiber-to-the-building, and fiber-to-the-business. These companies are leveraging many different protocols, including asynchronous transfer mode (ATM), Internet protocol (IP), and gigabit Ethernet.

In order to highlight the role of their active equipment, many equipment vendors represent the access network by showing the optical signal being generated in the central office or head end, passing through the PON cloud (that represents the outside plant), then somehow "magically" appearing at the customer premises (see *Figure 1*). From a carrier's perspective, it is imperative to realize that anywhere from 50 to 70 percent of the cost and a great deal of the complexity of deploying an optical access network are contained within the bounds of the "PON cloud."

Growing Bandwidth Demand

Over the past few years, it has been determined that the Internet has grown by as much as 100 percent every 100 days. Given this growth and the emergence of broadband applications, carriers are looking for ways to increase the capacity of their access networks and capitalize on incremental revenue-enhancement opportunities. The following dynamics and drivers are motivating carriers to consider optical access networks now more than ever:

- Carriers' desire to offer converged voice, video, and data services over a single network;
- Innovations and cost reductions in optical infrastructure and components;
- And evolving need for a protocol-ubiquitous, future-proofed access network.

At the same time, carriers will not deploy an optical access network simply to increase their network capacity or available bandwidth—it must enable the carrier to offer differentiable services and generate incremental revenue. Given the high cost of network installation and deployment, carriers do not want to deploy a broadband network with the anticipation of replacing it within several years; they want to build something that is going to enable them to be more competitive and increase their incremental revenue base over the long term.

The chart below depicts historical bandwidth demands, calculated by plotting the year that residential dial-up modem technology was introduced in the Americas. Given this development in dial-up technologies, America has experienced a historic annual increase in bandwidth of 60 percent per year.

If this historic trend is projected forward, it is anticipated that residential bandwidth demands will rapidly exceed 32 megabits (Mb) in the next decade. Given the networking technologies that are available (and in development), carriers will need to consider leveraging new technologies and deploying fiber deeper in their access networks to meet the ever-increasing demand for bandwidth and enable the delivery of revenue-generating broadband services.

PON Active Equipment Market

As PON technology continues to mature, many analysts and carriers alike are seeing tremendous opportunities for optical access networks to overcome the limitations of a copper-based access network and the traditional "bandwidth bottlenecks" that exist in last-mile networks. As the chart below depicts, CIBC World Markets forecasts tremendous growth in PON equipment sales over the next several years—with fiber-to-the-business representing the near-term growth opportunity and fiber-to-the-home representing more of a long-term growth area.

PON Architectural Considerations and Design Objectives

Carriers considering the deployment of a PON do not want to make that level of an investment only to be faced with the

FIGURE 1

The Emergence of the PON Cloud

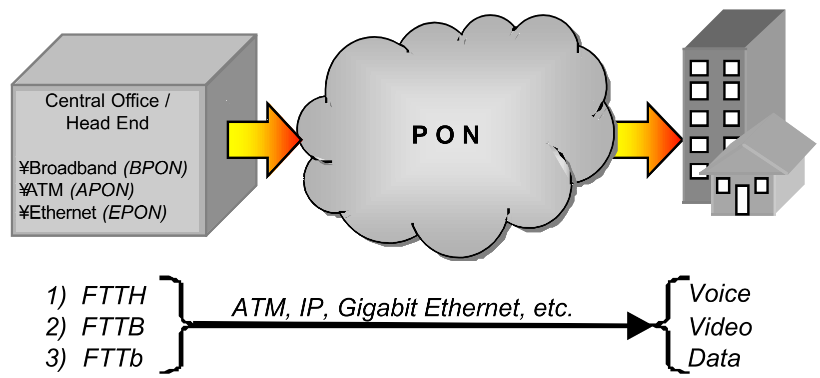

dilemma of redeploying or replacing segments of that network in the next several years in order to capitalize on the next generation of active equipment and broadband services. Carriers should plan and design their optical-access network deployments carefully to ensure that they create a network architecture that can be leveraged for the next 15 to 20 years—providing them with a ubiquitous approach to protocol and an infrastructure that will seamlessly support future technologies while minimizing initial capital expenditure requirements.

When designing a PON network, the carrier should carefully consider the advantages and disadvantages of the various PON architectures and ensure that their design includes or addresses the following objectives:

- Adaptability: ensuring that the physical plant will support different transport protocols, technologies, and emerging applications.
- Flexibility: enabling a highly configurable network through the deliberate placement of connectors, couplers, and splitters in strategic locations throughout the network design.
- Reliability: achieved by minimizing active equipment in the outside plant and limiting requirements for network powering and battery backup.
- Scalability: facilitating the incremental addition of subscribers over time and the ability to address new growth areas as they emerge.
- Low initial cost: enabling incremental infrastructure costs and deployment to be delayed until subscribers request service and/or revenue generation is ensured.

PON Architecture Designs

The generic architectural model (see *Figure 5*) below provides the general network layout and key infrastructure points in optical access networks. The network is generally comprised of three distinct network segments: feeder, distribution, and drop. The central switch is usually a central office or head end, which houses the active equipment (optical line terminal) for the optical access network. The local convergence point is the interface point between the feeder and distribution networks; whereas the network access point serves as the interface between the distribution and drop cable networks.

FIGURE 2

Growing Bandwidth Demand: Modem Technology

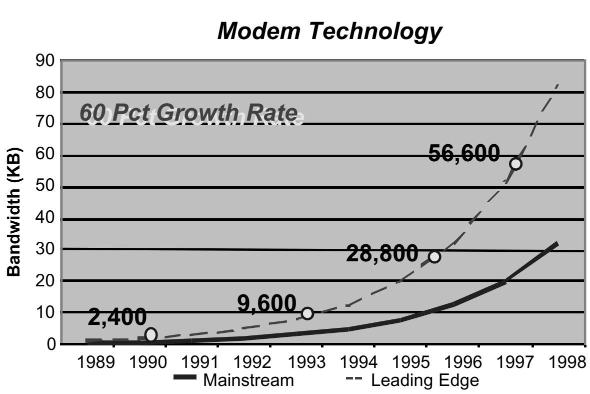

Source: TeleChoice

FIGURE 3

Growing Bandwidth Demand: Architectures

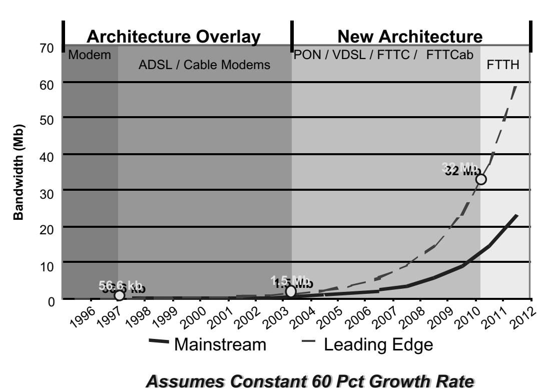

Source: TeleChoice

FIGURE 4

PON Active Equipment Market

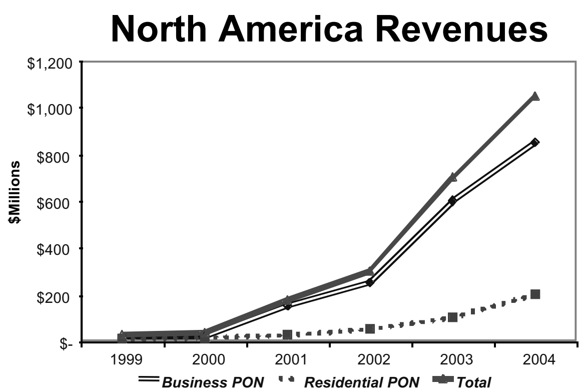

Source: CIBC World Markets

In general, there are three primary architectures being considered and/or deployed in access networks to date: central switch, local convergence, and distributed splitting.

Central Switch Architecture

The central switch architecture is the ultimate access network for every network operator—as the architecture provides for a dedicated fiber to connect each subscriber to the central switch (see *Figure 6*). There are numerous advantages to using the central switch architecture:

- It provides a dedicated optical path from the central switch to each subscriber,
- It provides the greatest upgrade potential with an almost unlimited bandwidth capacity,
- And it allows for subscriber configuration to be accomplished directly from the central office without a truck roll.

However, because of the fiber requirements of the central switch architecture, this architecture is very cost prohibitive—requiring extensive up-front capital expenditures that do not prove out for a majority of business cases. There are some applications where the central switch architecture is being used cost-effectively—mainly in Europe. In Europe, central offices are decentralized—making the central-switch location much smaller and closer to the end, and reducing the lengths of fiber required to reach the end subscribers.

Local Convergence Architecture

An emerging architecture—with a particularly strong potential in the United States—is the local convergence (LC) architecture (see *Figure 7*). LC leverages a single tier of couplers and a splitter to lean out the feeder segment and provided a dedicated optical path from the link control protocol (LCP) to each subscriber. The LCP is a very powerful point in the network because the carrier can actually run multiple PON extensions—with, for example, one extension providing service at a 1x8 split, while another is at a 1x32—and provide for extensive subscriber configurability. The LCP can also be a point in the access network that houses active equipment for a point-to-point network architecture (which does not include coupler/splitters).

The LC architecture provides carriers with the following benefits:

- It provides the best balance of capital expenditures and long-term scalability, and
- It eases subscriber configuration by reducing the number of network-interface points and providing a dedicated optical path from the LCP to each subscriber.

On the down side, if the LC architecture is deployed in a dense area, it may cause fiber density and footprint issues at the LCP.

Distributed Splitting Architecture

Up to 75 percent of most residential deployments today are deployed using the distributed-splitting architecture because of its low cost and lean fiber infrastructure (see *Figure 8*). The advantages of the distributed splitting architecture include the following:

FIGURE 5

Generic Layout

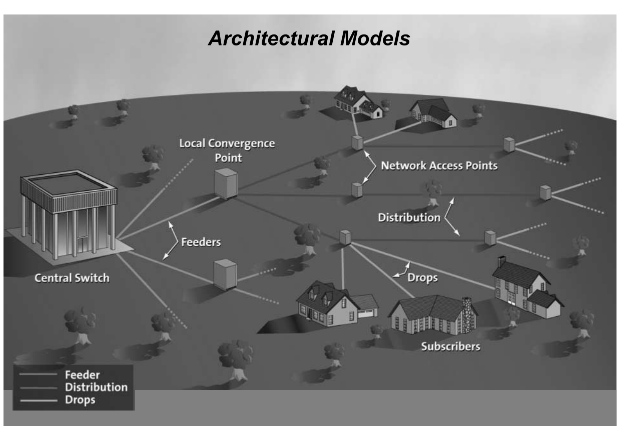

FIGURE 6

Central Switch Model

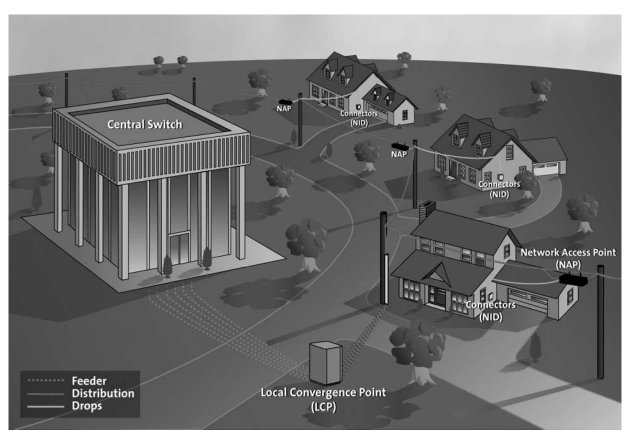

FIGURE 7

LC Model

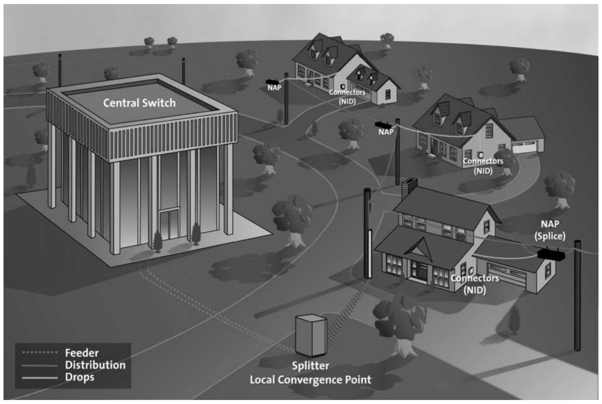

FIGURE 8

Distributive Splitting Model

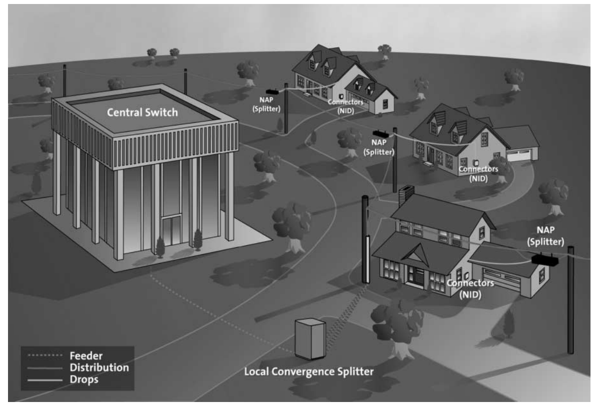

- It leverages the cumulative benefits of two tiers of couplers/splitters—one at the LCP and another at the network access point (NAP), and
- It is the least expensive PON architecture to deploy (for penetration rates > 30 percent).

However, the distributed splitting architecture may limit the long-term adaptability and scalability of the PON—particularly in its ability to support Lambda services. It is impossible to know today what active equipment or services will be available or required over the next 10 to 20 years. If at some point in the future the carrier decides it wants to migrate toward a dedicated optical path to each subscriber (LC or central-switch architectures), it literally has to recable its entire feeder and distribution-network segments—which would be extremely costly.

Network Migration Strategies

Anytime a carrier is deploying an access network, it is imperative that it develops and establishes a network-migration strategy. It must think past its current deployment and project the requirements of the next five to 10 years, expecting continual change and new technology, determining how to capitalize on these changes and increase revenues, and understanding what this will mean for its overall network design and architecture.

There are two of different ways to provide for network migration: physical infrastructure and active equipment. By focusing on migration through the physical infrastructure, a carrier can simply decrease its split ratio in the access network (1x32 to 1x16, 1x16 to 1x8, etc) or migrate from one architecture to another (distributed splitting to LC, LC to central switch) (see *Figure 9*).

By approaching network migration through active equipment, the carrier should plan to continuously upgrade its active equipment or purchase and install new equipment or remove any splitters that were installed at the central office or head end (see *Figure 10*).

Network Cost Modeling

When considering an optical access-network deployment, it is important to understand the various costs associated with

FIGURE 9

PON Cloud

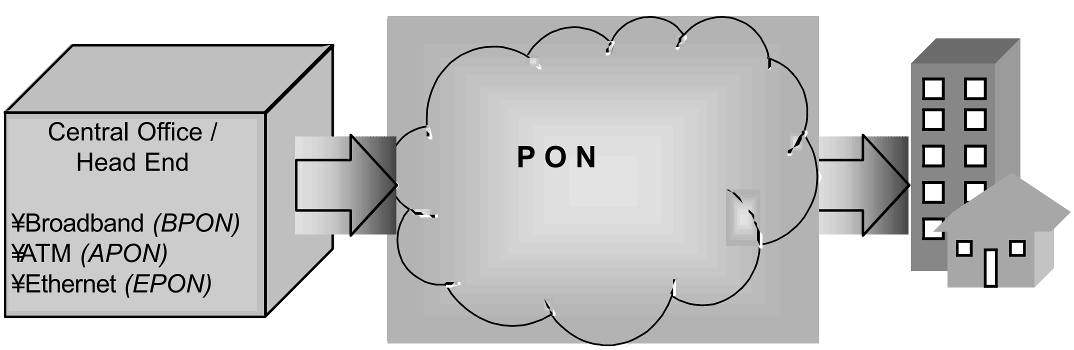

FIGURE 10

Active Equipment

1. *Decrease Splitter Ratio at Central Office/Head End*
2. *Upgrade Capacity of Existing Active Equipment*
3. *Purchase and Install New Active Equipment*

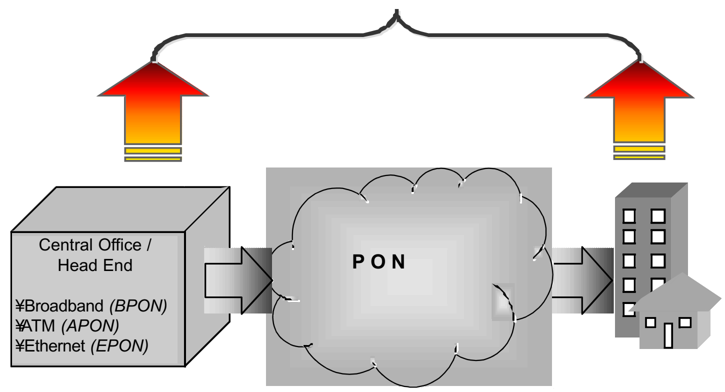

purchasing the passive infrastructure and how the infrastructure will be deployed and installed (see *Figure 11*).

A carrier must explore the immense number of deployment variables to gain a detailed understanding of how much a specific access-network deployment will cost (see *Figure 12*). Questions that address these network variables include *"How long is the feeder segment?,"* *"How many homes are going to be serviced?,"* and *"What is the anticipated take rate?"*

Labor and installation is a key piece of the deployment cost puzzle (see *Figure 13*). If a carrier deploys an optical access network that is complex or requires a high degree of craft skill, it is going to cost the carrier considerably more money than if installation complexity had been minimized. The carrier must be sure that it is looking at the deployment holistically and that it understands and is addressing each of the key variables towards minimizing the cost of deploying the network and reducing the return on investment period.

Depending on the deployment area and the associated business case, different PON architectures may be more advantageous in certain scenarios than in others. For example, many clients have gravitated toward distributed splitting as their PON architecture solution, believing it will be the least costly. In actuality—for certain scenarios—distributed splitting can actually prove to be more costly than LC architecture because of the high cost of couplers and splitters and potential for a low-port utilization in low-penetration areas (< 30 percent).

Conclusion

As the telecom industry continues to experience tremendous innovation in optical components and infrastructure,

FIGURE 11

Network Cost Modeling

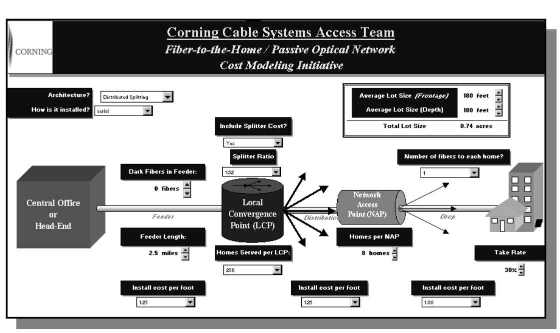

FIGURE 12

Access Deployment Cost Variables

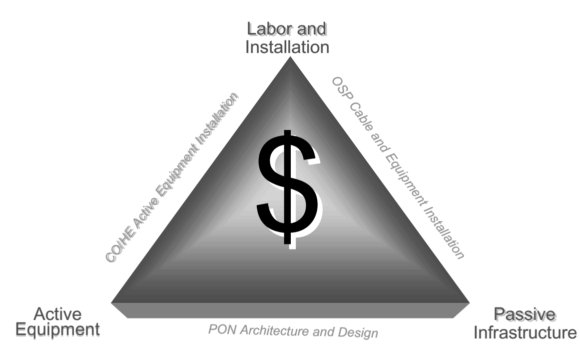

FIGURE 12

LCP Architecture

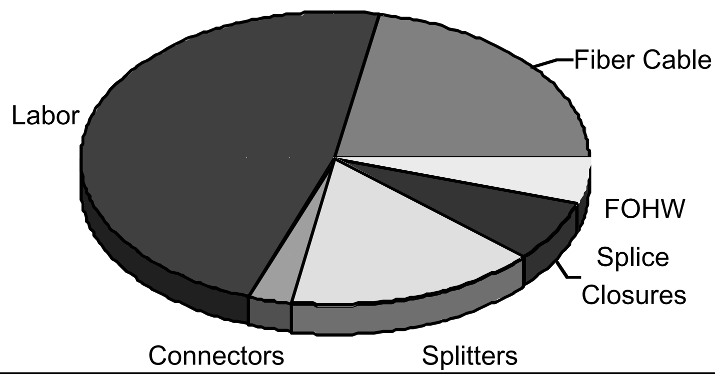

Relative Costs		
Fiber Cable	$1.00	22 Pct
FOHW	$0.23	5
Splice Closures	$0.30	7
Splitters	$0.73	16
Connectors	$0.13	3
Labor	$2.12	47

carriers are looking to deploy fiber deeper and deeper into the local loop; deliver differentiable, broadband services; and capitalize on incremental revenue enhancement opportunities. Prior to beginning *any* deployment, it is imperative that carriers gain a detailed understanding of the various cost drivers and dynamics that could positively or negatively impact the deployment of their optical access network and the resulting success or failure of their associated business case.

VoDSL for ILECs

Why RBOCs and IOCs Should Undertake In-Region Deployment of VoDSL Services

Steve Shaw
Director, Marketing
Jetstream Communications

Introduction

While the justification for the deployment of voice over digital subscriber line (VoDSL) services by competitive local-exchange carriers (CLECs) is clear to most in the industry, the same cannot be said for incumbent local-exchange carrier (ILEC) deployment of VoDSL. A number of ILECs, in fact, have only expressed the potential downsides of deploying VoDSL, such as the cannibalization of existing services. However, in-territory deployment of VoDSL can provide tremendous tactical and strategic benefits to both the marketing and operations organizations within ILECs. This white paper more closely examines the ILEC VoDSL opportunity.

For most CLECs, VoDSL represents the only profitably method for delivering a compelling bundle of voice and data services to the large and growing small-business and upscale residential markets. By enabling the delivery up 24 telephone lines and speed data services over a single leased copper off a single, centralized Class-5 switch, the economic drivers behind CLEC VoDSL deployment are quite clear.

Since ILECs already own the copper plant as well as Class-5 switches, or remotes, in every wire center, their motivations for deploying VoDSL services are not immediately apparent to many. However, upon closer examination, VoDSL can be shown to solve a number of near and long-term ILEC operational challenges. From a tactical perspective, VoDSL can immediately be leveraged to solve the growing copper exhaust challenge faced by many ILECs. From a strategic perspective, VoDSL provides ILECs with a graceful path for migrating their networks to a more efficient pure packet-switching architecture. From a marketing perspective, VoDSL can immediately begin to grow an ILECs average revenue generated per subscriber as well as provide a long-term method for mitigating the overall effect of competition.

VoDSL Dramatically Increases Operational Efficiency

By enabling ILECs to leverage their existing DSL broadband access networks and Class-5 local digital switches to deliver up to 24 telephone lines and a high-speed data service over a single copper pair, VoDSL can significantly improve the overall operational efficiency of ILEC local networks.

Single-Pipe Service Delivery

One of VoDSLs most significant contributions to operational efficiency is that it enables ILECs to maximize the value of every copper loop in their plant. Today, the vast majority of copper pairs are used to deliver a single analog telephone line. With VoDSL, an ILEC can now deliver up to 24 telephone lines and high-speed data services over a single copper loop. In other words, VoDSL enables ILECs to move to a single-pipe delivery method for all services delivered to subscribers (see *Figure 1*). Single-pipe service delivery can dramatically lower an ILECs overall capital and maintenance costs, and thereby significantly improve the profitability of all services delivered.

Today VoDSL Enables Single Pipe Delivery

More and more small businesses and consumers are ordering additional lines to accommodate growth and to receive new services such as high-speed Internet access. According to IDC, small businesses alone were expected to add 5.2 million new lines in 2000. While this is great news for ILECs, many of them are running into severe copper exhaust challenges and are being forced to run new copper loops to meet this demand. Unfortunately, the high capital and maintenance costs associated with running new copper can significantly impact the overall profitability of the service delivered. According to one major ILEC, they expected their second largest capital expense in 2000 to be maintenance of their copper plant. Running new copper pairs can also greatly increase the time it take to fulfill line orders, which can negatively impact an ILECs perceived service quality.

Since VoDSL can liberate more than 20 copper pairs per subscriber served, an ILEC can free up a tremendous amount of its copper plant by converting even a small percentage of subscribers to VoDSL service. With VoDSL, an ILEC can cap the size of its copper plant in existing markets and possibly reduce the overall size of its plant that must be maintained.

One other benefit of VoDSL is that it provides a key missing element for ILECs toward reaching of their goal of subscriber self-provisioning. Many ILECs can envision how to enable

FIGURE 1

The Single-Pipe Delivery Method

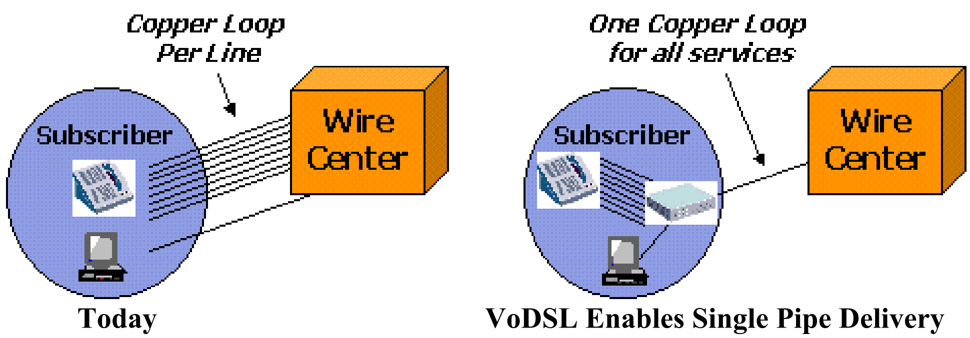

subscribers to activate various custom calling features on their existing lines, as it requires only software provisioning of the Class-5 switch. However, enabling subscribers to self-provision additional lines is much more problematic, as it requires physical activity at the wire center and often at the subscriber premises. With VoDSL, the activation of additional lines (i.e., integrated access device [IAD] ports) for a subscriber becomes solely a software-provisioning effort. By removing any physical provisioning required to activate an additional line, VoDSL takes ILECs one large step toward achieving their goal of subscriber self-service.

Centralized Switching

By leveraging ILECs existing digital subscriber line (DSL) broadband access networks, VoDSL enables them to begin the transition to a more efficient centralized Class-5 switching model. Instead of growing and maintaining the smaller Class-5 switches deployed today in every wire center (see *Figure* 2), VoDSL enables the use of fewer, larger, Class-5 switches housed in centralized regional switching centers (RSCs) (see *Figure* 3). These RSCs often serve today as tandem offices as well MegaPOPs that house the asynchronous transfer mode (ATM) and data switching equipment to which the digital subscriber line access multiplexers (DSLAMs) deployed in every wire center backhaul their traffic.

In addition to requiring fewer Class-5 switches, VoDSL also enable the use of "lineless" Class 5s. While the Class-5 switches used in today's network are filled with line cards that directly terminate copper pairs from subscribers on a line-by-line basis, the Class-5 used in VoDSL service do not require expensive line cards, since the VoDSL gateways that enable the service terminate directly into the concentrated GR-303 interface of a Class-5 switch.

Since VoDSL enables the use of fewer, and lineless, Class-5 switches, it can dramatically reduce an ILECs overall capital and maintenance costs. In addition, it enables the rapid introduction of new features, as there are fewer switches to upgrade or update with software releases.

A Logical Path to Pure Packet Switching

While most ILECs believe that the transition from circuit to pure packet switching is inevitable, there is much debate as

FIGURE 2

Today's Decentralized Switching Network

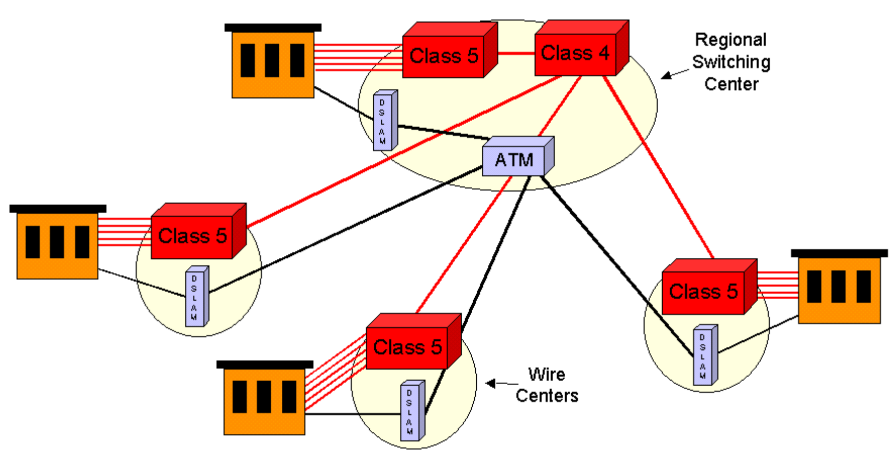

FIGURE 3

Centralized Switching Network Using VoDSL Broadband Access

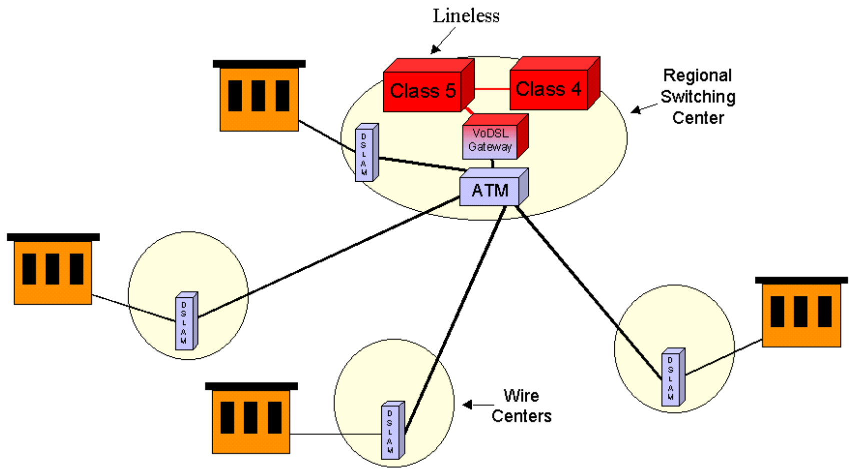

to when and how the transition will take place. In addition, there has been much discussion as to where VoDSL fits in the overall transition. However, it has recently become clear to many ILECs that VoDSL is directly on the path to soft switching. In fact, VoDSL is a necessary next step that must be taken in reaching that goal.

VoDSL enables ILECs to take the critical next step of using their broadband access networks to deliver voice services. VoDSL allows ILECs to continue to leverage the feature and reliability strengths of their existing Class-5 switches while immediately begin to realize the cost and service flexibility benefits of using their new DSL broadband access networks. Rather than a short diversion on the way to softswitching, VoDSL represents an absolutely necessary next step that ILECs must take in transitioning to a pure packet-switching model.

In a few years, when softswitches are better able to meet the feature and reliability expectations of the ILECs, they can be implemented in conjunction with the installed base of VoDSL gateways and IADs. The VoDSL gateways will be able to communicate directly with the softswitches and function as proxy servers that manage communications out to the IADs (see *Figure 4*). This hybrid model allows ILECs to continue to serve the majority of subscribers off of their existing Class-5 switches while slowing transitioning certain user groups over to more cost-effective softswitches.

VoDSL Provides Additional Near-Term Revenue

Some ILECs have viewed VoDSL solely as a means of overcoming operational challenges, such as copper exhaust. Other ILECs have viewed it as a method of mitigating the impact of competition by entrenching key

FIGURE 4

VoDSL Gateway Working with Softswitch

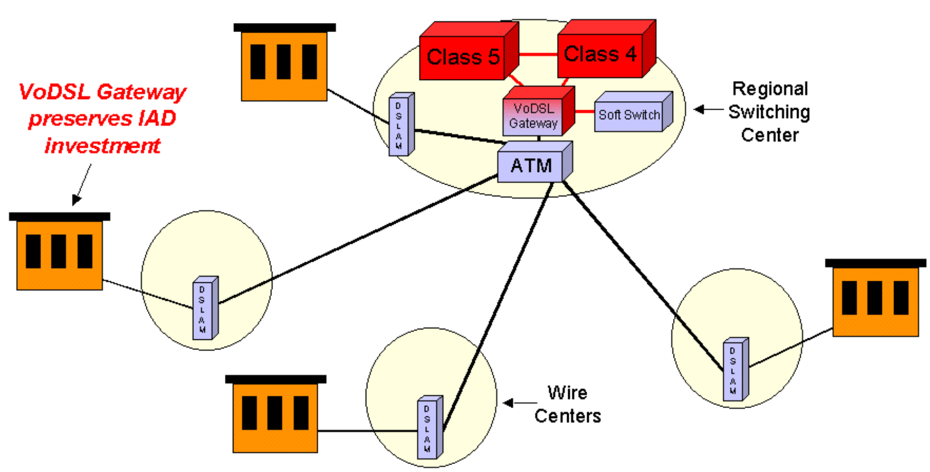

subscribers through service bundles, which is discussed further in a later section. However, VoDSL holds tremendous promise as a means of quickly growing the average revenue per subscriber by increasing the penetration of existing voice and data services as well as by enabling compelling new services.

Increases Penetration of Existing Voice and Data Services

VoDSL enables ILECs to further the success that they have experienced with service bundling. A number of ILECs have discovered that bundling enhanced features such as caller ID, call waiting, and voice-mail along with local service is a successful way to profitably increase their average revenue per subscriber. However, to date, service bundling has been primarily restricted to the delivery of various voice features. Recent end-user research has now shown a strong desire among target VoDSL subscribers to sign up for service bundles that include both voice and data services.

The theory behind service bundling is that an ILEC includes items in the bundle (in this case additional lines and features) at an effective price that is lower than the à la carte price. The subscriber is happy because they are getting services "cheap." The ILEC is happy because they are selling services that have a very low cost (much lower even than the discounted effective price) and that they would not have been able to sell (because the subscriber wouldn't have bothered to take them) if they were offered only at the regular retail price.

For example, say an ILEC subscriber currently has one phone line ($20) plus call waiting and caller ID ($3+$6). They receive their long-distance service from a different company, with which they are satisfied, as the subscriber believes they are getting a good deal at 7 cents per minute. They typically use 300 long-distance minutes per month. The subscriber by "default" would take a DSL service ($40) in addition to their current spending ($29). The total monthly revenue paid to the ILEC by the subscriber would be $69.

Now say the ILEC offers the subscriber a package for $129 per month (only available in package form) that includes DSL and three phone lines, all with call waiting, caller ID, and voice-mail, and 500 minutes of long distance. The indicated "value" of this package would be $180[1]. The ILEC's cost of the additional capabilities would be $20[2]. The ILEC has now taken their revenue up by $60 (with $40 of profit), so they are happy. And the subscriber is content, as they perceive they are saving $56. The subscriber has purchased more than they would have without the bundle.

Through targeted marketing of service bundles that include a specific number of telephone lines, an ILEC can significantly grow their overall line count. For example, targeting single-line DSL subscribers with a package that includes two feature-rich phone lines or targeting two- to three-line business subscribers with a bundle that includes four lines, or targeting five- to seven-line business subscribers with a bundle that includes of eight lines. The key point is that while such service bundles do have the potential to cannibalize some existing service revenue, there is much larger opportunity to grow total line count and service revenue through targeted selling.

Enables New Service Opportunities with Broadband Centrex

Small-business Centrex service continues to be one of the most significant unrealized ILEC revenue opportunities. According to Research First Corporation, there are more than 7.5 million small businesses in the United States, which, in total, utilize more than 38 million telephone handsets. However, more than 77 percent of these handsets (or 29.5 million) are served off of key telephone systems (KTSs). Since KTSs typically oversubscribe handset-to-analog trunks by more than 2:1, ILECs are providing approximately 14 million telephone lines to this subscriber group. If ILECs were able to go back into this installed base and successfully displace the KTSs with more reliable and feature-rich Centrex services, they could deliver approximately 15 million new telephone lines. In addition, while lines that service KTSs are typically feature poor, Centrex lines that service individual end users represent a significant additional source of revenue through the selling of custom local-area signaling services (CLASS), custom calling, and voice-mail features.

However, despite all its feature, growth, and reliability advantages, traditional Centrex service still faces a number of technical and business challenges that have prevented its widespread adoption within the small-business market. Largest among these hurdles has been the price of Centrex service. Due to the high costs associated with delivering traditional Centrex services, which requires a copper loop per telephone line as well as a lack of feature differentiation with standard business telephone service, ILECs have been unable to lower the price of Centrex service to the point where it is truly cost competitive with today's advanced KTSs and hybrid private branch exchanges (PBXs).

Fortunately, by using Jetstream's VoDSL solution to deliver Centrex services (a.k.a. Broadband Centrex), ILECs can overcomes these remaining hurdles and begin to successfully penetrate this 15 million new-line opportunity.

Since Jetstream's VoDSL solution enables ILECs to deliver up to 24 Centrex lines over a single copper pair, it dramatically reduces the cost of service delivery. Second is the unique and important feature of Jetstream's solution that enables ILECs to provide an oversubscribed Centrex service. With Broadband Centrex, a subscriber can now order telephone service using the same terms that they might use when buying a KTS, by specifying the required number of trunks and stations. For example, an ILEC could offer a 3x8 Centrex-over–DSL service, where eight Centrex lines would be provided, but only three lines could be simultaneously active. The eight lines delivered are standard Centrex direct inward dialing (DID) lines off the Class 5, supporting all available CLASS and custom calling features as well as meridian business sets, or P-Phones. The only, and key, difference is that the service is oversubscribed in manner similar to that of a KTS or PBX (see *Figure 5*).

By allowing for oversubscription, Jetstream's Broadband Centrex solution provides the key feature differentiation that now enables service providers to differentiate their Centrex pricing relative to standard business telephone service. In addition, since a Broadband Centrex service is delivered over the same DSL line, or pipe, as data service, it

FIGURE 5

Broadband Centrex Supports Line Oversubscription

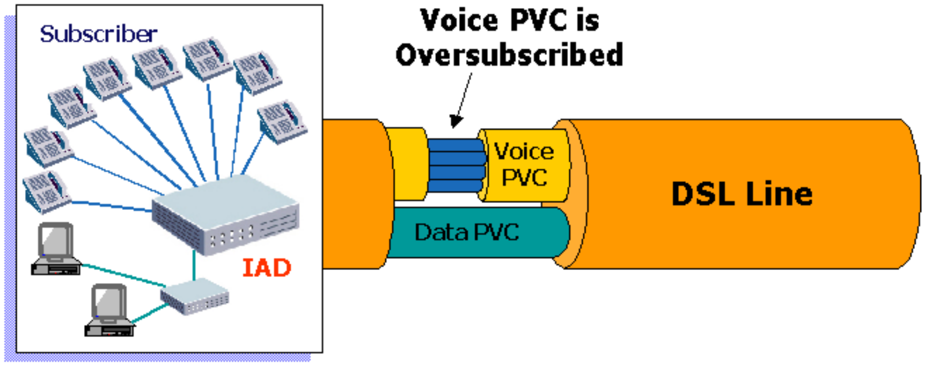

opens opportunities for voice and data service bundling, enabling further feature and price differentiation.

Another side benefit of Jetstream's Centrex oversubscription feature is that it can prevent the use of less expensive Centrex lines by ISPs looking to terminate a large number of dial-up connections.

VoDSL Mitigates the Impact of Competition

As competition continues to heat up in the local service arena, ILEC revenues are beginning to take a hit. CLECs have made tremendous progress in the few years since the Telecommunications Act of 1996. As of late 2000, there were more than 170 facilities-based CLECs competing in the United States, which, in total, generated annual service revenues of more than $39 billion. Together, they have more than 990 voice switches and 2,000 data switches installed. They have 9 percent of the access lines in the United States (more than 19 million lines), which is up from 3 percent in 1998. While the CLEC market has been experiencing some shakeout during the last six months, it continues to be a vibrant market and a real threat to ILECs.

Potential Loss Is Significant

While a number of early CLECs focused on serving the large enterprise market, the vast majority now targets small to mid-sized businesses as well as high-end consumers. While the monthly revenue contribution from individual subscribers in these markets may appear low, many CLECs have realized that the present value of each subscriber remains quite high ($200,000+ for a small business and $35,000+ for a high-end family) and that they constitute the largest and fastest-growing market opportunity for competitive service. In addition, these markets have traditionally been underdeveloped by the ILECs, and it has been relatively easy for the CLECs to get to the price-sensitive decision maker in these organizations. All in all, the CLECs represent a significant potential near-term and long-term revenue loss.

However, while the potential revenue loss may be large, the potential profit loss to ILECs is tremendous. Like all new competitive market entrants, the CLECs are first targeting the ILECs most attractive subscribers. In regulated markets such as telecom, this can have a profound effect on the overall profitability of the incumbent provider. This is especially true within the consumer market, where there are a large number of subsidized subscribers. In fact, one major ILEC indicated that less than 20 percent of their consumer subscribers contribute more than 100 percent of the overall profit in that market. And, it is exactly those 20 percent that the CLECs are targeting and into which they are beginning to make serious inroads.

VoDSL Entrenches Their Most Attractive Subscribers Today

VoDSL represents one of the best methods for ILECs to go back into their installed base and quickly entrench their most attractive subscribers. While CLECs are making tremendous progress, they have a long way to go before they have universal market coverage and the marketing resources by which to take full advantage of this opportunity. With VoDSL, ILECs can rapidly assemble, and profitably market, a compelling set of bundled services that would allow them to attract and retain their key small-business and consumer subscribers before the CLECs can cause more serious, long-term damage to their businesses.

Summary

Deployment of VoDSL represents a tremendous opportunity for ILECs. Not only does VoDSL solve a number of near-term tactical challenges, including copper exhaust and shrinking revenues per subscriber, but it also helps ILECs to solve a number of critical strategic challenges, including the difficult migration of their networks over to pure packet switching as well as the mitigation of the overall impact of competition to their businesses.

Notes

1. $40 + 3x$20 + 3x($3+$6+$6) + 500x$.07 = $180
2. Two VoDSL derived telephone lines with features: (2x$4) + (400 [actual long-distance minutes used] x $.03) = $20

Synchronizing Deep Fiber Baseband Access Network Design with Traditional HFC Infrastructure

Donald Sorenson
Manager, Advanced Access Planning, Transmission Network Systems Sector
Scientific-Atlanta, Inc.

Abstract

Increasingly, optical network technologies are being evaluated for their suitability as a residential access network platform. Clearly optical networking holds promise for the distant future as networked applications evolve to demand greater transport capabilities. However, today and for years to come, hybrid fiber/coax (HFC)–based access systems, with their low-cost structures and evolving performances, will remain the dominant access network of choice for delivery of interactive multimedia services in the majority of the served residential markets. In a contemporary context, it is quite likely that an all-optical access platform could be best utilized as a strategic tool tailored to delivering high-value business class services to 10 percent or less of a residential serving area and a much larger percentage of the small to medium-sized businesses presently passed by HFC feeder fiber. A low first-cost optical access network may be an ideal strategic/offensive overlay to an existing or new HFC network.

This paper explores the key applications engineering issues associated with such an overlay and proposes a methodology for synchronizing key optical access and HFC network elements. The paper concludes with a detailed analysis of how optical split ratios and cable sheath fiber counts can impact plant first costs. This work is focused on the engineering issues associated with layer 1, the physical layer, of the seven-layer open systems interconnection (OSI) network model. Commentary on upper-layer requirements and protocols are limited to issues that impact the logistics of an overlay deployment and subscriber provisioning.

HFC and FTTx Architectures

The overlay of contemporary deep fiber architecture on an existing or new HFC plant is facilitated by effectively coordinating the location and functionality of key network elements. A brief review of each access method's legacy and contemporary structural fundamentals will serve as a useful point of reference as various design and cost sensitivities are explored later in the analysis to follow.

Hybrid Fiber/Coax

Routinely, the architectural details of HFC plants are varied to meet a host of constraints associated with factors such as geography, home density, upgrade logistics, and costs, as well as many other details. However, fundamentally most plants adopt a variant of a ring(s)—star(s)—bus configuration, as illustrated in *Figure 1*. Typically, one or more fiber rings interconnect primary head-end and hub facilities, in some cases secondary rings link head ends with secondary hubs or optical termination nodes (OTNs) that commonly serve tens of thousands of subscribers. In turn, fiber nodes are typically linked to hubs, OTNs, or head ends in a star configuration.

Ring and meshed configurations offer physical path diversity and associated enhancements in plant reliability. As high-performance service and business models stimulate the need to extend fiber deeper into the plant, a greater use of rings is anticipated at deeper levels. Within this trend, network functionality is also on the move—network process functions are migrating to the edge of the plant where they can be better tailored and incrementally scaled to meet subscriber and new business demands. A third ring extending from a secondary hub through a series of smaller serving areas to remote star locations at the "1,000 to 4,000 homes passed" level effectively enables the deployment of advanced video and cable data services [1] and is a logistical requirement for the deployment of an all-fiber access overlay, as will be shown later in this paper.

Optical Access Network Overlay

Various all-optical access networks are increasingly seen as the following:

- The enabler of high-value services to a large under-served population of small business and home offices

FIGURE 1

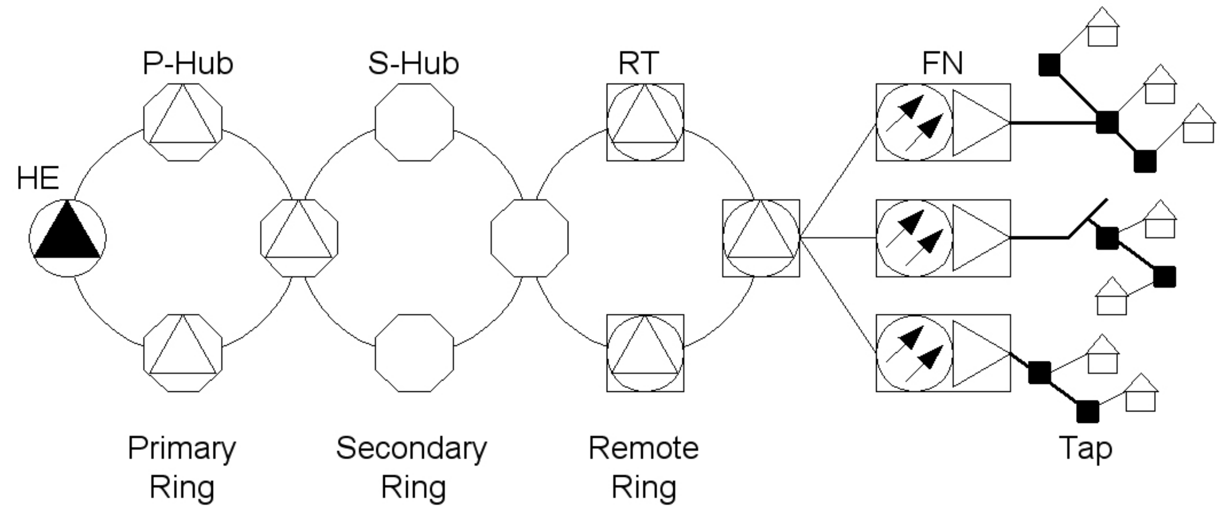

- A low-risk infrastructure enhancement that has the potential to extend a plant's useful life well beyond that of copper-based technologies

A contemporary all-optical access architecture must be optimized to cost-efficiently deliver advanced services to a small percentage of a wide subscriber market base being passed.

To insure the integrity of these services and an extended useful plant life, the network must be reliable and easily scaled with increasing take rates without major upgrade or modification.

By contrast to most copper-based networks, all-optical access networks do not utilize active amplification or signal-processing elements in the transmission path and are referred to as passive optical networks or PONs. Ideally, the active elements at both the subscriber and service-provider ends of the network are safely housed in conditioned facilities, and only passive components such as fiber cable and splicing devices suffer the environmental rigors of the outside plant, as shown in *Figure 2*.

Rarely would a dedicated fiber connection, between the service provider's internal network and a residential or small-business subscriber, be cost-effective. To distribute costs a means of allowing multiple subscribers to access a network, fiber is implemented through the use of a multi-port optical splitter or coupler (see *Figure 3a*).

FIGURE 2

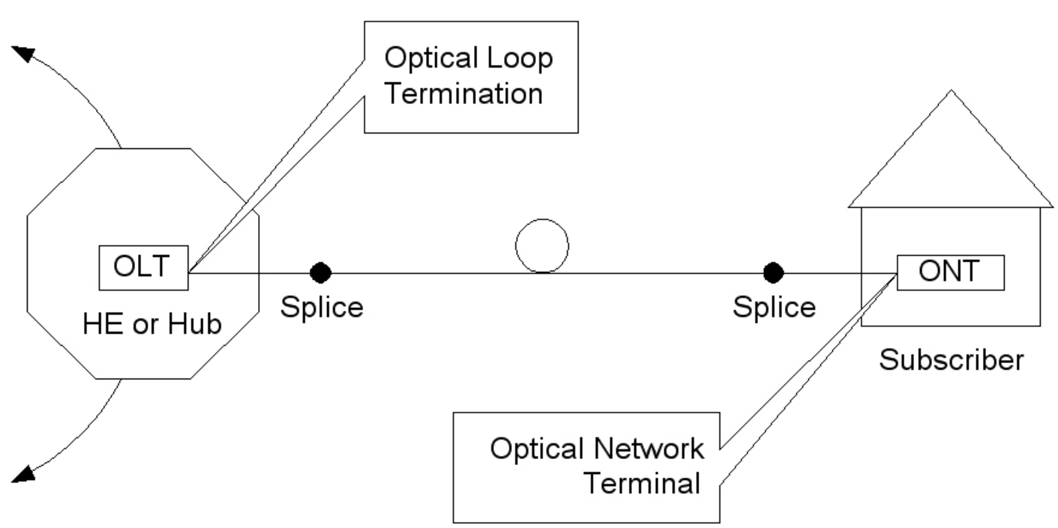

FIGURE 3A

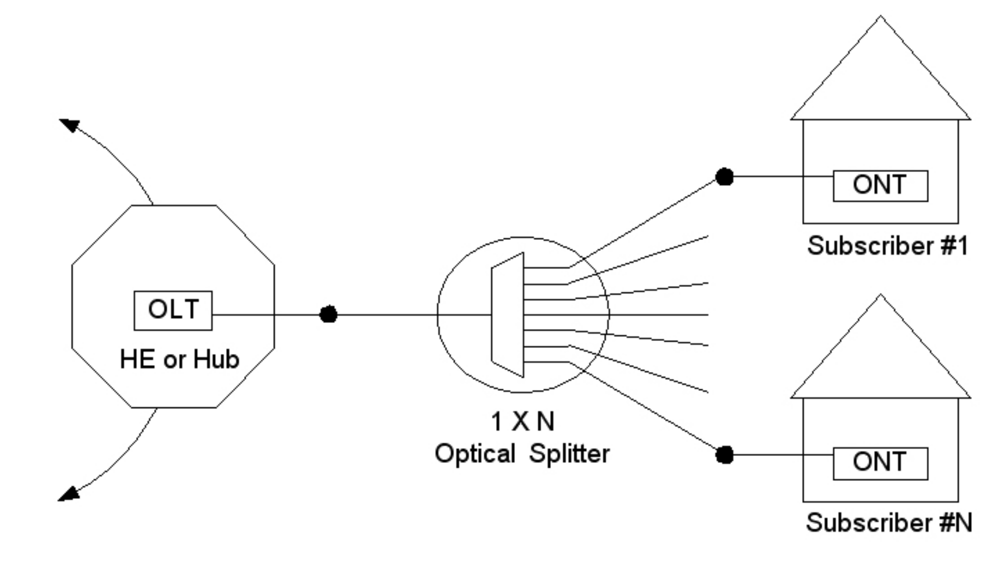

FIGURE 3B

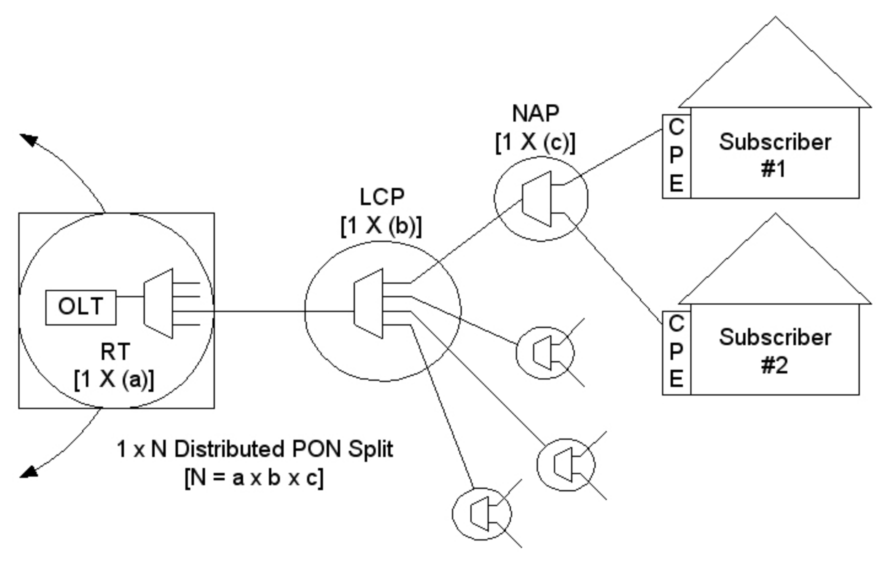

In this instance, each PON fiber becomes a shared media. If the PON is used in conjunction with dense wavelength division multiplexing (DWDM) or course wavelength division multiplexing (CWDM), each subscriber is allocated a dedicated wavelength, or subchannel carrier on a wavelength, for their dedicated use. Using a time division multiplexing (TDM) approach, the downstream traffic is broadcast to all of the subscribers served by a splitter, and upstream access is controlled via a suitable media access control (MAC) protocol, a number of which are competing for industry acceptance. For reasons relating to cost and operational logistics, it is often not desirable to lump the splitting function into one physical location [2,3]. As shown in *Figure 3b*, the splitting function can be divided between multiple locations along the fiber path. The product of the split ratios at each location becomes the PON's overall return fiber to subscriber ratio.

The sizing and distribution of the PON splitter function is the principal tool used to synchronize an optical access network's elements with an existing or new HFC plant.

As previously mentioned, ideally a PON's active elements are housed in conditioned environments. However, the realities of applying an optical network design to a wide serving area typically demand that active electronics be placed in the outside plant. At the subscriber side of the network, the logistical realities of service activation place the ONT [4] outdoors on the exterior of the subscriber premises as customer-premises equipment (CPE). On the network side, construction costs and operational issues typically demand a high degree of fiber aggregation deep in the design. Without such measures, fiber counts quickly become difficult to manage with even the highest of PON split ratios. Splicing costs escalate dramatically while network reliability and scalability suffer. Typically, it is not practical to provide path redundancy on one or a few PON fibers; however, when hundreds of fibers collect on their way to a network facility, the failure group size and recovery time from a fiber cut become unacceptable. Furthermore, high fiber-count head-end or hub runs are very likely to exceed existing fiber inventories and subsequently adversely impact plant first costs.

The location of an optical loop termination (OLT) [4] synchronizes quite well with the evolution of the third-ring nodes or remote terminals (RTs) [1] in advanced HFC architectures.

Commonly, an OLT or RT location will service a 1,000 to 4,000 home area. Construction practices demand the consolidation of fiber cable sheaths as they route to the RT. These local consolidation points (LCPs) [2] (see *Figure 3b*) provide another physical location to easily distribute a portion of the network's total split ratio.

The proper allocation of PON split ratio to the LCP can be used to synchronize its physical location in the network with that of each physical HFC fiber node, thus providing a convenient fiber path for the HFC node and common physical point of equipment access.

Access to the PON fibers passing a subscriber premises is achieved in manner similar to an HFC coax tap—the principal difference being that the equivalent PON tap is fed by a dedicated PON fiber that is not shared by other upstream or downstream PON taps (see *Figure 3b*). To logically clarify this difference, the PON fiber tap point is referred to as a network access point or NAP [2]. Fundamentally, a NAP is composed of a suitably sized ground or aerial mounted fiber splice enclosure near a small group of subscribers. The optical service drop to the subscriber is terminated at the CPE using an optical splice or connector. At the NAP, the drop will be spliced either directly on one of the PON fibers passing through the NAP or onto a port of an optical splitter located in the NAP.

The installed first cost and service-activation cost of a PON are quite sensitive to the fraction of the total PON split ratio allocated in the NAP. In general, a higher split ratio at the NAP yields a lower first cost of construction and slightly higher cost per subscriber served should the take rate at the NAP exceed 50 percent of its capacity.

Another challenging issue regarding PON architectures is that of remote active element powering. Historically, remote network device power has been supplied via the metallic signaling media. Without such a media, a PON must either include some form of copper-based powering network, or its remote actives must be locally or subscriber-powered. Both alternatives have their advantages and difficulties, a selection between which will largely be determined by business and public-policy issues.

The issue of powering remote PON equipment in an HFC overlay is facilitated by utilizing the powering signal commonly available from the underlying cable television (CATV) distribution coax.

By including a light gage twisted pair with the subscriber's optical drop network, power can be provided to the CPE using industry-standard power passing tap devices. In those cases where spare power capacity is not available from the HFC network, other twisted pair powering alternatives exist that can be implemented at the time of plant construction at a modest percentage of total PON overlay cost.

Network Design Constraints

The practical objective of a design is the implementation of a serviceable network that meets near- and long-term needs for the lowest possible cost. Compliance with this objective in legacy copper-based technologies is strongly enforced by well-understood and refined engineering practices. In part, the scope of these practices must be expanded when new technologies, such as PON, are adopted. The following material outlines key engineering, construction, and operational considerations that will act to constrain the design variables associated with implementing a PON overlay.

Serving as a reference point for the following discussion, *Figure 4* is a simplified illustration of a PON overlaying an HFC fiber node with a physical size of 256 homes-passed (HP). In turn, the physical node (PN) is fed by a 1,000 HP logical node (LN) or RT.

Network Engineering Considerations

Routinely, more than one communications service provider serves a residential or commercial market area. This implies a reduced initial service activation rate as market-share is obtained. Accordingly, new plant designs of any type are commonly engineered for a lowest possible first cost that enables the network to physically address the market area. This condition is particularly true with regard to PON deployment. It is anticipated that early take rates for PON–based services will be quite low—less than 10 percent of HP. This condition is likely to remain until the value of service portfolios, exclusively deliverable via an optical network, increase and construction costs associated with PONs decrease.

With a lowest possible first cost, a PON overlay must address 100 percent of a market service area, and with minimal provisioning or manipulation. the network can be made available to any one of the subscribers passed.

In keeping with a low first cost and take-rate assumptions it is anticipated that PON split ratios will be under-subscribed by as much as 50 percent and that initial RT data-processing capacity will be over-subscribed by as much as 10:1. These two factors substantially lower first costs with regard to OLT ports, associated packet switching and fiber counts, and/or advanced optical wavelength provisioning in the upper network transport rings. A 50 percent under-subscription of a PON split ratio implies that less than 50 percent of a PON's available optical splitter ports will be terminated and in service at any one time. Dedicated wavelength division multiplexing (WDM)–based PONs inherently provide this ability. However, when using a shared access TDM PON control method, this engineering option is enabled by the selection of a MAC protocol that has the ability to dynamically allocate access to the PON over a specified range of possible subscriber counts. This capability greatly simplifies provisioning and record keeping while allowing the network to provide bandwidth to customers based upon the terms of their subscription agreements without regard to the number of subscribers terminated on a given PON. Conversely, and based upon optical budget limits, which as

FIGURE 4

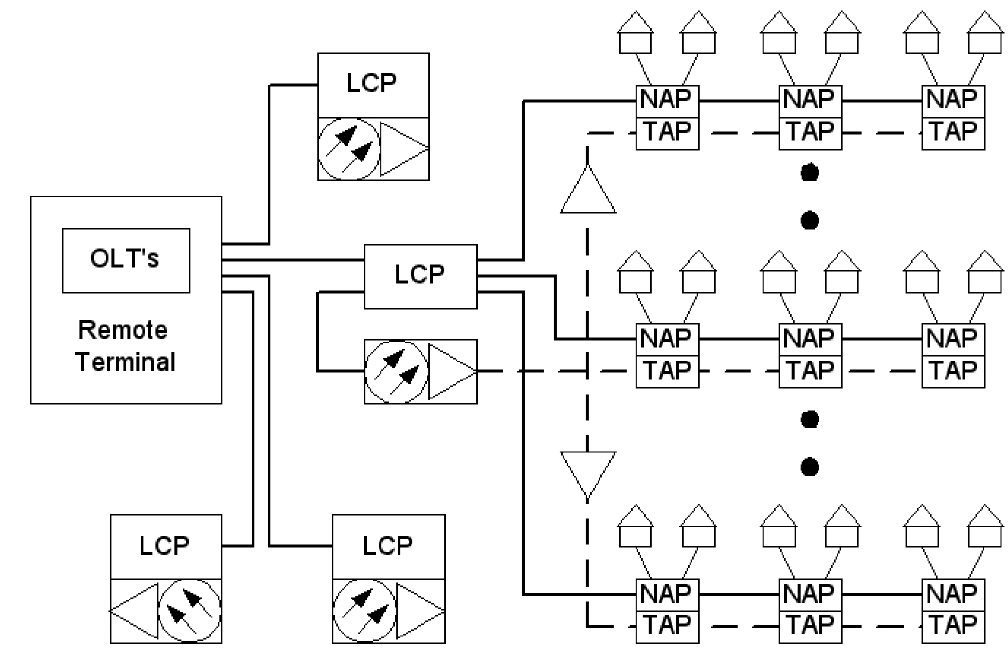

discussed later, under-subscription allows the network engineer to pass lit fiber by all of the homes in a given serving area without excessive OLT costs.

The use of WDM or the selection of the proper TDM MAC protocol and the planned under-subscription of a multi-port PON allows the network to distribute lit fiber past 100 percent of the subscriber base without encumbering excessive OLT provisioning cost or degrading subscriber data throughput.

If the PON is to be under-subscribed, it is advisable to take measures that help ensure the under-subscription through the life of the network. This can be achieved by distributing a portion of the PON split ratio to the LCP. By doing so, each of the PON's ports can be spread across multiple distribution branches or NAP locations that are not geographically adjacent. Routinely, subscription rates in small clusters of homes passed can reach 100 percent. By spreading a PON's ports across multiple distribution branches, the likelihood of 100 percent utilization is diminished.

The PON must be capable of supporting the eventual migration of standard services delivered today via copper media. This requirement principally impacts optical budget limits that in turn drive PON split ratios and passive device performance limits. Large PON split ratios, such as 1:32, demand optical physical interface (PHY) transmitter power levels of up to +20 dbm. Elevated levels such as this are achieved at a cost premium and bandwidth penalty over more common PHY transmit levels of +10 dbm or less.

In the years to come, the need to operate the PON at multi-gigabit speeds is almost certain. Large PON split ratios and the associated elevated PHY transmit power levels are likely to delay the point in time where an existing network can be easily upgraded to multi-gigabit operation.

Thus PON splits such as 1:8 or 1:16 are attractive from the point of keeping CPE PHY–device costs down and ensuring early migration to higher data rates. It may also be advantageous to allocate some portion of the PON split ratio to the OLT side of the network. For example, a four-port directional coupler can facilitate insertion of overlaying ls to support additional services in the future.

Construction Engineering

Minimizing construction first costs are key to the financial success of a PON overlay. Low first cost enables early economic business models that rely on the delivery of high-value services to a selected few of the subscribers passed by the network. Unfortunately, only a small number of PONs have been deployed, and related experiential knowledge is limited at best. The following points address construction issues impacted directly by PON architecture decisions.

Through the analysis to follow, it has been found that PON construction costs are reasonably insensitive to actual fiber counts. However, costs are particularly sensitive to the labor associated with fiber management, termination, and optical splitters, along with related outside plant enclosures and places to locate them. Again, the size and distribution of PON split ratios can be optimized to minimize these initial costs.

Analysis has shown that poorly designed 1:32 or 1:16 PON can be significantly more expensive than an optical network based on a dedicated home-run fiber for every home passed!

PON fiber counts at a single location in the network can easily escalate to levels that are exceedingly difficult and costly to manage. Additionally, in the event of a fiber cut, a large count represents a correspondingly large failure group size. For example, 200 – 1:16 PON fibers under-subscribed by 50 percent still represents 1,600 high-value paying subscribers that will lose service when a pole falls or a cable is dug up. With cable restoration time in mind, it is advisable to keep individual fiber sheath counts below 50 if possible and fiber aggregation point (LCP and RT) counts well below 150. This can be achieved by adjusting the amount of PON split ratio performed at the NAP and LCP.

Increasing NAP split ratios reduce fiber counts at the LCP and move the LCP into the network, which can be used to synchronize the LCP location with an existing or planned HFC fiber node. Correspondingly higher LCP split ratios have the same effect on the RT.

A uniform bill of materials composed of industry-standard elements that are frequently reapplied throughout the construction footprint enforces cost control and reduces construction-schedule delay risks. The selection of one standard cable sheath fiber count for all distribution branch runs and another for feeder runs between the LCP and RT is beneficial. The variables involved in selecting a distribution branch fiber count are NAP split ratio(s), anticipated home density range, HFC distribution branch length (typically on a per active basis), and an allocation for spare fiber. *Table 1* illustrated how NAP split ratios can be selected based upon home density and branch length. In this case, higher NAP port counts are implemented using two smaller optical splitters. Such an approach allows the deferral of the optical splitter and its associated cost.

A generous allocation of spare fiber throughout the design will insure restoration and application flexibility in the future and possibly aid in addressing legislative issues associated with equal access. Adding individual fibers to sheath counts has a relatively minor impact on first cost. Later, these fibers can be used to deliver high-value multi-gigabit services via dedicated port connections from facilities remote to the subscriber, such as a primary head end or hub.

At the time of this paper's authorship, the industry does not provide a selection of low-cost standard NAP enclosures suitable for co-location with standard and power-passing cable television taps. This condition will change as PON deployments become more routine.

Operations Engineering

Easy PON subscriber provisioning and low or no PON equipment maintenance are key operations cost-control objectives. To meet these, objective measures must be taken at the physical layer of the PON as well as at the MAC protocol layer if TDM is used.

Installation of a subscriber service drop and CPE must be executable by a technician with a relatively low craft skill level and without the need of advanced tooling or diagnostic equipment.

Turn-up of the CPE must be a plug-and-play event, with the MAC protocol automatically handling provisioning and network access issues in much the same way as done today by cable modems via the data over cable service interface specifications (DOCSIS) protocol. Dedicated WDM approaches will be challenging in this regard, due to the inherent need to manually manage or allocated wavelengths and WDM filters to individual subscribers. For this reason, WDM solutions are likely more appropriate for high-bandwidth services to small and medium-sized businesses.

Optical budgets, noise, and return loss margins must allow the use of industry-standard mechanical fiber splice devices at both ends of the service drop. Experience has

TABLE 1

Distribution Branch Cable Sheath Fiber Count Scaling

Length	900' Typical		
Homes Passed	8 HP	16 HP	32 HP
CATV Tap Size	2 Port	4 Port	8 Port
Taps / Branch	4	4	4
NAP Split	1:1	2x1:2	2x1:4
PON Fibers	8	8	8
Spare Fibers	4	4	4
Total Fibers	12	12	12

shown [3] that cable drop damage due to bending or dirty connections can occur during the installation process. The ability to easily open connections for optical time-domain reflectometer (OTDR) measurements at the NAP and below the last optical splitter is convenient. Additionally, the PHY receivers at both ends of the PON must provide a dynamic input range that can tolerate the initial exclusion of the NAP's optical splitter.

The ability of the PON to tolerate the exclusion of the NAP splitter for the first NAP subscriber is more difficult to achieve when high optical split ratios are applied at the NAP.

The service technician cannot be required to access or open adjacent NAP or LCP enclosures to establish to which PON fiber the subscriber's NAP has been allocated nor to activate or light the PON fiber. Having to access other devices during the provisioning process can be logistically difficult and is certain to significantly drive labor costs up on a per-home-served basis. Again, this objective is directly influenced by the allocations of splitter ratios at the NAP and LCP, as well as cable sheath fiber count.

The LCP must be affordably provisioned such that each distribution cable sheath has at least one dedicated fiber for each NAP on the sheath.

This fact coupled with a desire to keep first costs down encourages higher branch cable sheath fiber counts and lower LCP split ratios. During plant construction, the appropriate fiber color for each individual NAP is tagged with a color marker on or in the NAP, or alternatively the appropriate fiber can be ring-cut from the sheath and stored ready for use.

PON Overlay Synchronization

The proposed methodology of synchronizing a PON overlay with an existing or new HFC plant is driven by the objective of having the PON's key functional elements closely located, both physically and logically, with the corresponding HFC devices. Practically, direct physical co-location is required at the CPE, tap, and RT levels. Strict co-location of the LCP with a HFC fiber node may be desirable but is not a requirement. The accomplishment of this objective must, of course, align with as many of the previously mentioned engineering constraints as possible.

The PON design is driven by three key parameters: the PON split ratio, the ratio's distribution through the network, and the sheath/enclosure fiber count limits. These three parameters narrow the field of alternatives such that final selections can be made based upon specific performance or functionality objectives and first-cost reduction measures.

PON Split Ratio Analysis
All possible ratio distributions between the NAP, LCP, and RT were explored for the combined split ratios of 1:1, 1:2, 1:4, 1:8, 1:16, and 1:32. The first three ratios offer limited flexibility with regard to distribution but were explored for comparative purposes. The first column of *Table 2* lists the possible ratio combinations for each overall split ratio.

Limiting the associated fiber counts to a maximum of 136 controls the LCP's position or depth in the network. The LCP fiber counts shown in *Table 2* are inclusive of both inbound and outbound fibers. With regard to split ratio allocation, the LCP's position in the network is dominated by the NAP. For example, a 1:4 NAP ratio favors a 256 HP LCP position.

TABLE 2

Split Combo	LCP / Physical Node		Rermote Terminal		Normalized	
NAP-LCP-RT	Size	Fibers	Size	Fibers	First Cost	Ultimate Cost
1-32 PON Split Ratio						
1-32-1	128 HP	132 F	4096 HP	128 F	2.3	2.3
1-16-2	128 HP	136 F	2048 HP	128 F	2.3	2.3
1-8-4	64 HP	72 F	1024 HP	128 F	1.7	1.7
1-4-8	64 HP	80 F	512 HP	128 F	1.7	1.7
1-2-16	64 HP	96 F	256 HP	128 F	1.6	1.6
1-1-32	64 HP	128 F	128 HP	128 F	1.7	1.7
2-16-1	256 HP	136 F	4096 HP	128 F	1.9	2.0
2-8-2	128 HP	72 F	2048 HP	128 F	1.2	1.4
2-4-4	128 HP	80 F	1024 HP	128 F	1.3	1.4
2-2-8	128 HP	96 F	512 HP	128 F	1.2	1.4
2-1-16	128 HP	128 F	256 HP	128 F	1.2	1.4
4-8-1	256 HP	72 F	4096 HP	128 F	1.0	1.3
4-4-2	256 HP	80 F	2048 HP	128 F	1.0	1.3
4-2-4	256 HP	96 F	1024 HP	128 F	1.0	1.3
4-1-8	256 HP	128 F	512 HP	128 F	1.0	1.3
8-4-1	512 HP	80 F	4096 HP	128 F	0.9	1.2
8-2-2	512 HP	96 F	2048 HP	128 F	0.9	1.2
8-1-4	512 HP	128 F	1024 HP	128 F	0.9	1.3
1-16 PON Split Ratio						
1-16-1	128 HP	136 F	2048 HP	128 F	2.3	2.3
1-8-2	64 HP	72 F	1024 HP	128 F	1.7	1.7
1-4-4	64 HP	80 F	512 HP	128 F	1.7	1.7
1-2-8	64 HP	96 F	256 HP	128 F	1.6	1.6
1-1-16	64 HP	128 F	128 HP	128 F	1.8	1.8
2-8-1	128 HP	72 F	2048 HP	128 F	1.2	1.4
2-4-2	128 HP	80 F	1024 HP	128 F	1.3	1.4
2-2-4	128 HP	96 F	512 HP	128 F	1.2	1.4
2-1-8	128 HP	128 F	256 HP	128 F	1.2	1.4
4-4-1	256 HP	80 F	2048 HP	128 F	1.0	1.3
4-2-2	256 HP	96 F	1024 HP	128 F	1.0	1.3
4-1-4	256 HP	128 F	512 HP	128 F	1.0	1.3
8-2-1	512 HP	96 F	2048 HP	128 F	0.9	1.2
8-1-2	512 HP	128 F	1024 HP	128 F	0.9	1.3
1-8 PON Split Ratio						
1-8-1	64 HP	72 F	1024 HP	128 F	1.7	1.7
1-4-2	64 HP	80 F	512 HP	128 F	1.7	1.7
1-2-4	64 HP	96 F	256 HP	128 F	1.7	1.7
1-1-8	64 HP	128 F	128 HP	128 F	1.8	1.8
2-4-1	128 HP	80 F	1024 HP	128 F	1.3	1.4
2-2-2	128 HP	96 F	512 HP	128 F	1.2	1.4
2-1-4	128 HP	128 F	256 HP	128 F	1.3	1.4
4-2-1	256 HP	96 F	1024 HP	128 F	1.1	1.4
4-1-2	256 HP	128 F	512 HP	128 F	1.0	1.3
8-1-1	512 HP	128 F	1024 HP	128 F	1.0	1.3
1-4 PON Split Ratio						
1-4-1	64 HP	80 F	512 HP	128 F	1.7	1.7
1-2-2	64 HP	96 F	256 HP	128 F	1.7	1.7
1-1-4	64 HP	128 F	128 HP	128 F	1.9	1.9
2-2-1	128 HP	96 F	512 HP	128 F	1.3	1.5
2-1-2	128 HP	128 F	256 HP	128 F	1.3	1.5
4-1-1	256 HP	128 F	512 HP	128 F	1.1	1.4
1-2 PON Split Ratio						
1-2-1	64 HP	96 F	256 HP	128 F	1.8	1.8
1-1-2	64 HP	128 F	128 HP	128 F	1.9	1.9
2-1-1	128 HP	128 F	256 HP	128 F	1.4	1.6
1-1 PON Split Ratio						
1-1-1	64 HP	128 F	128 HP	128 F	2.1	2.1

Alternatives can also be explored, such as creating a four-port NAP utilizing two two-port splitters at the NAP; this will favor a 128 HP position at the expense of doubling the number of fibers in each cable branch. At first, such an alternative may be dismissed as costly, but individual fibers are relatively inexpensive compared to optical splitters. Having two fibers available at each NAP further delays the purchase of NAP splitters until three or more NAP ports are in service, which may never happen. The measure also reduces the dynamic range requirements imposed on the CPE and OLT PHY receivers.

Strategies that postpone the use of a splitter at the NAP significantly impact the average subscriber provisioning costs, particularly during the early years of plant operation while subscription rates are low.

The RT is positioned in the network by following a process similar to that used with the LCP. The total incoming fiber count is held constant by adjusting the number of LCPs reporting to the RT. As shown in the fifth column of *Table 2*, the total inbound LCP fiber count was held at 128. This quantity was chosen, rather than the previously mentioned 150, to allow for outbound or ring fibers and HFC node fibers that were not included in the analysis.

Both NAP and LCP ratios influence the depth or position of the RT in the network. NAP/LCP ratios of 1:4/1:2 favor a 1,000 HP RT in the 8, 16, and 32 PON splits.

Physical placement of ground mounted 1,000 and 2,000 HP RT cabinets is common, particularly with regard to digital loop carrier (DLC) telephone systems. Small, 500 HP RTs can be pole mounted, but become numerous and pose powering challenges when extended backup times are required to ensure network operation during extreme weather conditions.

The cost data provided in *Table 2* is based on items that generate a relative cost difference between PON split options. For example, the fiber cost for a PON split of 1 – 1 – 32 would be considerably greater than 32 –1 – 1; this along with other associated impacts such as number of RTs required for each combination. The resulting dollars per home passed were then arbitrarily normalized to the 1:4 NAP option. Normalization of the data is appropriate since the cost analysis is gauging relative merits rather than absolute values.

The cost impacts of various subscription take rates are difficult to predict and are not included in the analysis. This is principally due to the cost of an optical splitter at the NAP if it is needed at the time of service provisioning. In many instances, a splitter may not be needed for the early subscriber(s) whose drops can be connected directly to the PON fiber allocated to the NAP. Thus, unless there is a great deal of subscriber clustering, the impact of NAP splitter costs will be negligible until take rates approach 30 percent to 50 percent. Due to this variability, the analysis focuses on the initial relative cost and that of a 100 percent take rate. The 100 percent figure reflects the comparative cost between low and high subscription-rate pockets.

The principal conclusion drawn from the cost data is that from a relative-economics point of view, a plant can be designed by appropriately applying a wide range of split combinations.

This is good news from a plant-engineering point of view, providing the network designer latitude to meet the numerous other non–cost-related engineering constraints such as those previously outlined.

Observations

The key take-away observations from this analysis effort and discussion are as follows:

- Properly executed, PON optical split ratios do not greatly affect network first costs.
- The appropriate allocation of PON split ratios between key network elements can be effectively used to control the physical and logical location of those elements in the network.
- Control of sheath fiber counts can be effectively used to defer optical splitter cost at the NAP and to ease performance demands on PHY receivers and transmitters.
- One or two standard sheath fiber count configurations can be readily adapted to varying subscriber densities and HFC tap configurations.
- Mechanical fiber splice connections at the NAP and CPE are required to keep provisioning costs down and support drop troubleshooting. The network's optical performance requirements must be tailored accordingly.
- The eventual migration of the network to gigabit speeds is likely to occur earlier with lower overall PON split ratios, particularly with regard to TDM–based solutions.
- The planned under-subscription of a PON split ratio can be leveraged to reduce the cost of the RT while enabling lit fiber to pass by 100 percent of the serving area's subscriber base.
- The utilization of a self-provisioning MAC protocol capable of dynamic bandwidth allocation without regard to PON split or subscription is required to maintain high customer satisfaction and low provisioning/maintenance cost.

This analysis was focused on PON split ratios in exponential powers of 2. Clearly, other ratios are equally valid and have been used successfully [3], such as multiples of 3. A significant cost benefit would not be anticipated in alternative ratios; however, they may offer useful alternatives with regard to LCP and RT placement and provisioning.

Summary

Fundamentally, HFC alone is a proven financial workhorse capable of delivering outstanding residential video, voice, and data services. Based upon the work done thus far, it is clear that it will be some time before the PON alone can compete economically in a standard residential market space with HFC. However, in part, this fact is contrasted by three emerging near-term trends:

- The growing utilization of e-business practices by small, under-served businesses—businesses within easy service-provisioning range of an HFC plant.
- The steady demand for advanced data networking and voice services for work-at-home and small branch office employees.
- A strong financial desire to reduce or eliminate the risk of obsolescence in newly constructed plant. Extending

the plant's operation during its positive cash-flow life by as little as two or three years has significant investment repercussions.

Perhaps the optimal migration compromise between the realities of today's economics and the future's service demands may be the strategic application of optical network overlaying a new or existing HFC plant.

Accordingly, this paper has proposed a series of engineering considerations and design methodology for the synchronization of a PON with a modern HFC network. The methodology has been demonstrated analytically to be capable of effectively placing key PON elements at physical and logical locations consistent with an underlying CATV network.

References

[1] P. Connolly, "New Carrier Class Architecture Busts Bandwidth Blues," NCTA Technical Papers (2001).

[2] M. Conner and E. Dunton, "Future Proofing Your Residential Broadband Access Network," Technical Proceedings, 2000 National Fiber Optic Engineering Conference (NFOEC).

[3] G. Mahony, "Fiber to the Home: A Summary of Bellsouth's First Office Application," Technical Proceedings, 2000 National Fiber Optic Engineering Conference (NFOEC).

[4] "Broadband Optical Access Systems Based on Passive Optical Networks (PON)," ITU–T Recommendation G.983.1.

Challenges and Opportunities in Home Networking

Matt Tormollen
Vice President, Marketing
BroadJump, Inc.

Introduction

Home networking has emerged as the leading application for residential broadband, with more than one-third of all current subscribers using some type of home network. The proliferation of inexpensive, plug-and-play network address translation (NAT) routers has given the early adopters of broadband an easy solution for sharing a high-speed connection among multiple machines. Broadband service providers are keenly aware of this trend but have encountered a number of challenges in capturing the value of home networking, which has thus far gone mainly to networking hardware vendors and consumer electronics retailers.

Service providers have, however, absorbed many incremental costs associated with home networking. Provider networks have experienced an increase in average bandwidth utilization, and support-desk phone calls are increasing in volume and length because of the added complexity of troubleshooting connection problems with a NAT router in place. The home-networking story for service providers has largely been one of increasing costs without incremental revenue—not exactly the business model they are looking to achieve.

Beyond this frustration, however, lies a new wave of opportunity. Home networking devices have fallen enough in price that hardware vendors are now looking to bundle data over cable service interface specifications (DOCSIS) and digital subscriber line (DSL) modems into their boxes and begin selling directly to service providers. In addition, less experienced broadband subscribers are more receptive to paying an incremental monthly service charge in exchange for supplying the equipment and supporting the installation and maintenance of a home network. Finally, the profit potential of advanced hardware functionality, such as voice over Internet protocol (VoIP) and home monitoring, can only be realized in conjunction with back-end network hardware and systems. Service providers are now in a position where, armed with the right tools, they can leverage advanced home networking capabilities to create new revenue streams and transform what has been only a source of additional costs into a profit generator.

First Steps in Home Networking

Service providers looking to offer a home networking solution to their subscribers will face a number of difficult decisions. Will they serve primarily as hardware resellers, looking to capture a one-time margin on home networking equipment? Or will they offer home networking as a service, giving subscribers the option of paying an incremental monthly service fee in exchange for free equipment and ongoing support? Finally, what technology solution will they focus on, be it traditional Ethernet, wireless, or some other alternative, such as HomePNA? But regardless of the business model that service providers choose to pursue, the first avenue of opportunity lies in the reduction and limitation of costs associated with NAT home networks. The cost drivers that providers need to focus on first are the expense of helping a subscriber successfully set up a home network, and then allowing the subscriber to self-manage and troubleshoot their network on an ongoing basis.

Automating Installation

Usability studies show that the average broadband subscriber encounters a number of issues when setting up even the most basic form of NAT–based home networks. Ensuring that each machine being connected to the network is qualified for broadband service, setting up each machine for dynamic host configuration protocol (DHCP), and then booting up the NAT device and every machine in the correct order are just a few a the frequent headaches that generate calls to the help desk. Just as service providers need to ensure that all of the computers at different subscriber locations are configured properly to receive service in the wide-area network (WAN), each subscriber needs a means through which to standardize their multiple machines for service in the local-area network (LAN). The answer to this problem lies in giving subscribers intelligent software tools that automate the qualification and installation process. This will help them to quickly configure each machine for broadband Internet access, step them through the appropriate process of connecting their machines to a NAT hub, and then instruct them on how to restart the devices in the correct order to power up the network to a working state.

Home Network Support
While the point of installation represents a large component of cost in offering a home-networking program, subscriber support cannot end there. Once a NAT home network is installed correctly, service inquiries related to the NAT device itself should be minimal, primarily because these devices are typically reliable and simple in their operation. The problem, however, is that the presence of a NAT device makes diagnosing and resolving problems with computer configuration and network connectivity far more difficult and complex. Support-call duration is often greatly extended because a representative must ask a series of probing questions to determine the root cause of problems in the subscriber environment. In addition, simple tools, such as a PING test, are unusable in a situation where the subscriber's computer is invisible behind a NAT router. Providers often have no record of the machines that have been installed in a home network, which further extends the process of gathering the necessary information. To cost-effectively manage the support of residential home networks, providers need to deploy an application that allows subscribers to self-diagnose potential problems within their network and, if they are unsuccessful, allows customer-service representatives (CSRs) to quickly upload information on each machine and its status.

Home Networking As a Service

The broadband providers that will achieve the greatest success in deriving new revenue streams from home networking components will view these hardware devices as the enabler of new, paid-for services. The opportunity to transform NAT functionality into a revenue-generating service is now emerging—primarily because the subscribers that are just now adopting broadband service are less sophisticated in managing their own customer-premises equipment (CPE) and will pay an incremental fee for provider support. And beyond the revenue potential of NAT lies the next generation of services that advanced home networking devices will help to create. Linking multiple types of devices, adding secondary VoIP phone lines, and offering home monitoring and home-security functions will all depend on a centrally located gateway device in the home that the service provider can interact with and maintain control over. Today's home networking devices give subscribers the ability to change configurations and activate new services, but they do not rely on head-end systems to coordinate the delivery of new, Internet protocol (IP)–based services. Ultimately, service providers will look to combine home networking hardware with intelligent software applications that make these incremental services available to subscribers in a secure, reliable, and paid-for format.

Promoting and Selling Home Networking
The most fundamental aspect of transforming home networking into a service is the ability to sell the service to a new subscriber at the appropriate touch point. Once subscribers have already installed their services, demand for a home networking device will likely lead them to a retail outlet, rather than back to their broadband service provider. For this reason, providers must be able to anticipate the demand for home networking services and make these offers intelligently and efficiently during initial contact with a new subscriber. By allowing subscribers to upgrade their service plan with home networking equipment, and then helping them to install their home network quickly and easily, service providers will establish themselves as the first source of both networking hardware and associated services.

Silent Functionality
To create multiple revenue-generating services through a piece of home networking hardware, service providers must be willing to subsidize the deployment of devices with embedded functionality over and above their baseline capability. This "silent functionality" becomes the enabler of services for which subscribers will pay incrementally on top of their basic access plan. To allow for this process to take place, however, the provider must deploy a layer of software intelligence that presents these advanced capabilities to the subscriber in the form of service offerings and allows the subscriber to activate the services by agreeing to the associated surcharge. In addition, the software interface with which the subscriber interacts must be able to communicate with back-end systems, coordinating the adoption of new services with the provisioning of network elements that work with home networking CPE to deliver services to the subscriber in the manner specified.

Integrated Billing
The creation of new, paid for services through home networking necessitates that providers have the ability to efficiently and cost-effectively bill for incremental offerings. Just as subscribers will view home networking hardware components as services, they will expect to be billed for these services in a logical and accurate manner. To ensure that this happens, providers must have intelligent software within their network that registers the activation of a new, paid-for service and initiates a corresponding event within the provider's billing system. This is made possible only by a client-server architecture that gives subscribers control over service selection, while at the same time, a layer of abstraction from the critical back-end systems must register and bill for these services.

Home Networking Solutions

Broadband service providers looking to transform home networking hardware into profitable, revenue-generating services will have several key decisions to make in defining their approach to this opportunity. At the most basic level, providers will need to select a business model. Proposed home networking models vary widely in terms of upfront investment and potential return, so providers will ultimately have to determine how strategic home networking is in reference to their long-term objectives. Second, any home networking program will involve the deployment of home networking hardware, but the control layer can be either hardware or software-based, so providers will need to define a network topology that maximizes their overall investment. Finally, providers must consider the various services that they want to leverage from home-networking capabilities. Advanced gateways promise to deliver a host of advanced services to the residential broadband subscriber, but each category of service involves significant investment in the face of uncertain market potential.

Choosing a Business Model
Given the multitude of potential models that exist for deriving value from home networking, it would seem that identifying one particular program with a high degree of success

would be next to impossible. Service providers have only to look at their basic access priorities, however, to select a generic model that holds out the most promise. First, a service provider will want the payback period of a home networking service to be as short as possible, and second, the provider will want to align revenues with expenses. Since these factors eliminate the possibility of either selling equipment at retail prices (a one-time gain with long-term costs) or fully subsidizing the cost of all hardware components in exchange for a monthly fee (an unacceptably long payback period), it is logical that the best approach is a hybrid model.

As a hypothetical example, consider a home-networking program that involves a choice of three CPE devices: a simple modem, an integrated Ethernet router/hub and modem, and an integrated 802.11b wireless router/hub and modem. In this program, subscribers choose one of these three devices at the time of ordering broadband service and pay incremental monthly fees for more advanced functionality. The service provider offers network interface cards (NICs) at retail price and supports these cards accordingly. By adopting this model, the service provider earns enough additional revenue on networking CPE devices to cover their incremental cost in a matter of months, and they earn a margin on the sale of each NIC. Subscribers, in turn, derive value from this model in that the incremental monthly charge saves them the cost of paying for the router/hub up front (which would cost them at retail much more than the price at which the provider can purchase these devices), and they pay the same price for NICs that now come with comprehensive support at installation and on an ongoing basis. In this model, the service provider has a modest upfront investment on each subscriber, creates a sustainable stream of service revenue, and sets the stage for future cash flows derived from advanced services.

Control—Hardware or Software?

To create and sustain profits from home networking, the service provider must have a control mechanism through which services can be provisioned, managed, billed for, and finally, deactivated. While every home-networking program will necessitate the deployment of an integrated system, the provider can allocate different levels of control between the hardware and software elements within the network. Hardware-based control, in the form of integrated circuit (IC) functionality, has the benefit of being stable and preintegrated, while software-based control, in the form of embedded agents with server-side applications, has the benefit of being more cost-effective and remotely updateable. The key issue for a service provider, however, is that a hardware-centric approach requires a single source solution, which in turn exposes the provider to a number of additional risks from a pricing and technology advancement standpoint. For this reason alone, the service provider will find that intelligent software provides both advanced functionality and control, along with the flexibility to manage several hardware platforms, in a model that is economically advantageous.

Enhanced Services

Without question, the first application of home networking is NAT. While NAT holds the promise of significant revenue streams, the true long-term potential for home networking lies in services that go beyond personal computer (PC)–based data transmission. With this potential for reward comes risk, and service providers will need to carefully evaluate the services that they are willing to subsidize in the short-term to create an installed base of devices waiting to be unlocked for the delivery of enhanced, next-generation services. VoIP, home monitoring, and security, as well as gaming and entertainment services, are all potential add-ons that providers can build into their home-networking program. What services will justify the cost? Many industry reports, especially when analyzed collectively, would seem to indicate that all of these services hold enormous revenue potential and that service providers should focus on all of them, lest they miss out on the true "killer app" for broadband. What service providers need above all else, however, is a home networking platform that is flexible enough to deploy these services cost-effectively yet can be updated with more advanced capabilities should a robust market for any of these particular services start to emerge. With this goal in mind, service providers will find that a minimal amount of added hardware components, in conjunction with a software platform that can manage the remote update and activation of services, will give them this type of flexibility.

Summary

Broadband service providers are experts at deploying new systems in their network to enable the delivery of additional, paid-for services. With the onset of home networking, a host of new services can be made available to subscribers in this same manner. Service providers will be constrained by the economic realities that they faced in deploying their traditional service offerings and accordingly will need to make investment decisions in home-networking programs with the same level of consideration. The right business model, combined with the appropriate combination hardware elements and software intelligence, will equip service providers with the means to fully capitalize on the emerging home-networking market opportunity.

Scalability of the Optical Core Network

Lord Uva

Product Manager, Photonics Business Unit

Cisco Systems

This paper will discuss the need for scaling the optical core, as well as the scalability trend in general. It will also talk about scalability of the optical core in terms of capacity, reach, and speed (CRS). Finally, it will review some of the emerging optical technologies, at a very high level, that are enabling the scalability of optical networks.

Scalability Drivers

What is driving scalability? A couple of years ago, if a manager needed to get in touch with an employee after hours, he had to wait until the next day. These days, he can send him an e-mail at 2 am and expect him to respond. The Internet has changed everything—the way we work, live, play, and learn. Now people are working 24 hours a day. They cannot just work eight hours and go home and relax until the next day. People are sending e-mails to them and sending their files across continents. We have seen the emergence of Internet protocol (IP) applications in voice and wireless, as well as the continued expansion of e-business. Essentially, we are experiencing an exponential increase in Internet traffic. All of this is putting increased pressure on the core of the optical network. In other words, to be able to accommodate all these changes, the core of the optical network needs to be very scalable. It needs to be able to scale in channel capacity, reach, and speed and at the same time lower the cost per transmitted gigabit.

Figure 1 is a look at pre-Internet traffic patterns. Before the Internet, enterprise networking was essentially a local-area phenomenon in which about 80 percent of the traffic was maintained within an enterprise campus or building. Only about 20 percent of the traffic actually left the enterprise environment, as applications such as e-mails and file transfers. The challenges of the pre-Internet era were mainly scaling local access network (LAN)-type equipment such as Ethernet and fiber distributed data interface (FDDI) switches.

With the Internet traffic pattern, the 80/20 percent rule still applies. The difference is that now 80 percent of the traffic actually leaves the enterprise campus or building and only 20 percent stays within. This means that there is more traffic coming out of enterprise environments and onto the optical core or backbone of long-haul networks. This is one of the reasons the optical core needs to be scaled to meet growing bandwidth demand.

The Internet has changed the world. The role of a service provider (SP) is to provide a seamless extension of Internet applications, not just within the enterprise but beyond it as well. A scalable infrastructure is needed to support all of this. In the Internet economy there are suppliers, partners, and customers all over the globe trying to conduct business at any time of the day. SPs have to find ways to build scalable networks to handle this unpredictable Internet traffic crossing their networks.

There are essentially three ways to scale an optical network. First, the optical core can be scaled by increasing the channel capacity of optical systems. This is accomplished by utilizing tighter channel spacing in the conventional (C)-band and/or migrating to long (L)-band or, in the future, to short (S)-band. Second, the optical core can be scaled by upgrading to higher bit rates through faster transponders. Third, optical-core networks can be scaled by increasing span distance or reach. However, there is always a tradeoff between speed, reach, and capacity. These issues are constrained by amplifier bandwidth and optical power budgets.

Scalability Trends

What is the biggest challenge for the optical core network? How can it be scaled? *Figure 2* shows the scalability trends in the core optical network. A few years ago, when dense wavelength division multiplexing (DWDM) first came out, people were talking about systems with 200 Gigahertz (GHz) channel spacing. Right now, there are systems out there with 100 GHz and perhaps 50 GHz channel spacing. We have seen DWDM systems go from 2.5 gigabits per second (Gbps) to 10 Gbps in the last two years, and there are 40-Gbps systems being developed in various labs across North America. With the preponderance of data/IP traffic in carrier networks, we now see systems moving from long haul to ultra-long haul (ULH), the latter being those systems with longer span reach than conventional systems that require regeneration, reshaping, and retiming (3R) every 500km or so. This reduces equipment cost by more than 30

FIGURE 1

Pre-Internet Traffic Pattern

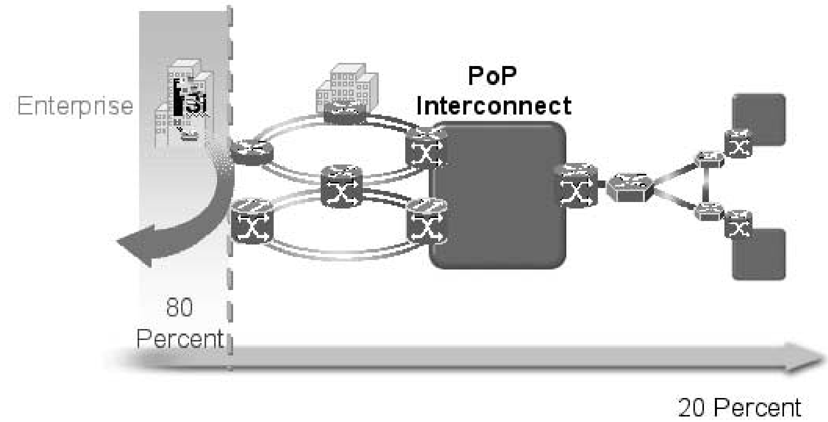

percent. *Figure 3* shows the projected market for ULH; as you can see, it is projected to be about $3.6 billion in 2004.

Basically, all these trends are being driven by the use of the Internet, as the SPs need to be able to scale their networks to accommodate bandwidth demand brought about by exponential increase in Internet traffic.

Technologies

The emerging technologies for scaling optical networks are being made possible by smart engineering and a better understanding of physics, with most of the innovation taking place at the component level. For example, Raman amplifiers have now made it possible to increase the channel capacity of networks by tapping into L-band and, eventually, S-band.

Some companies are using dispersion-managed solutions to go longer distances, and there has also been an increase in the use of out-of-band forward error correction (FEC) transponders for extended long reach. In the next couple of months, we should have available better-performance, high-speed modulators for 40-Gbps systems. Right now, there may not be any commercially deployed 40-Gbps out there, and one of the problems is that polarization mode dispersion (PMD) is a very big issue for 40-Gbps systems, especially for old fiber types. Because of the stochastic nature of differential group delay, dynamic PMD compensators will be needed.

In summary, the Internet is changing everything. There is a lot of pressure on the core of the network, and the next-generation Internet will require a scalable optical core.

FIGURE 2

Scalability Trends in Optical Core

Channel/bit rate migration

8	16	32	40	64	80	128
400 GHz	200 GHz	100 GHz	50 GHz	25 GHz	12.5 GHz	
2.5 Gbps	10 Gbps	40 Gbps				

FIGURE 3

ULH Market Forecast

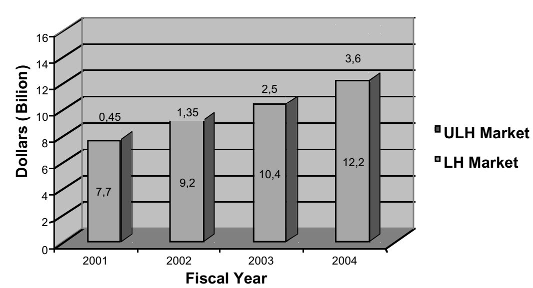

Data from RHK - StarTrax 2000 - November 2000

The Evolution of Wavelength Switches

Ian Viney
Senior Director, Product Marketing
Luxcore

Introduction

Dense wavelength division multiplexing (DWDM) networking is evolving at a phenomenal pace, particularly in the area of wavelength switching. This paper examines the evolution of wavelength switches, comparing today's first-generation architectures with soon-to-be-released second- and third-generation alternatives. This information will assist service providers in making more informed network decisions, thereby avoiding stranded capital and loss of their competitive edge.

A New Unit of Capacity

Since the first deployment of DWDM technology in the mid-90s, the capacity of optical fiber has increased by a factor of nearly 200, with no signs of a slowdown. In addition to enabling this astounding bandwidth growth, DWDM also introduced the wavelength (or lambda) as a new unit of network capacity. As shown in *Figure 1*, situated between the multiple-terabit capacity of a single fiber and the relatively small 50 megabit-per-second (Mbps) building block of the synchronous optical network (SONET) hierarchy, the wavelength is a perfect addition to the grooming hierarchy.

Unfortunately, today's switching equipment is unable to groom at the wavelength level, and so operators are forced to use traditional SONET cross-connect equipment as an expensive stopgap. As a result, the development of an all-optical switch capable of cost-effectively grooming signals at the wavelength level has become a major thrust in the telecommunications industry. This paper examines the evolution of wavelength switching and compares alternative architectures, both in service and being proposed in the marketplace.

The Ideal Wavelength Switch

Before comparing alternative switch architectures, it is worth defining the "ideal" wavelength switch as a basis for evaluation. In this paper, the term "wavelength switch" refers to a device that terminates multiple optical fibers, each carrying multiple wavelengths (see *Figure 2*).

From a grooming perspective, a wavelength switch can transfer the entire contents of any wavelength on any incoming fiber to any wavelength on any outgoing fiber. Wavelength connections are established on the order of tens of milliseconds under the direction of provisioning staff.

Wavelength switches do not groom traffic at the sub-wavelength level (e.g. synchronous transport signal [STS]–1). Instead, STS–level traffic are aggregated onto dedicated fibers and passed to a lower-level SONET switch to be groomed.

In many respects, the high-level requirements of a wavelength switch are identical to those of an electronic switch—non-blocking, scalable, automated, cost-effective, etc. In addition, however, DWDM has opened up an exciting new possibility—transparent wavelength switching.

A transparent wavelength switch grooms individual wavelengths carried by a DWDM fiber without imposing restrictions on either the rate or the format of signals carried on those wavelengths. The benefits of transparent routing are enormous. For example, signal rates can be upgraded (optical carrier [OC]–48 to OC–192 to OC–768), and new protocols deployed (e.g., 10 Gigabit Ethernet [GbE]), without any hardware changes to the switch. In contrast, optical-electrical-optical (O–E–O) switches are not rate or protocol independent (referred to as "opaque") and require a forklift upgrade to achieve the same results.

In an era when subscribers expect service changes and upgrades to be instantaneous, and in an era when service-provider competition is at the highest level ever, many of the benefits of instant-rate and protocol change are obvious. Maybe not so obvious, though, is that transparent wavelength switching also enables a wide range of new offerings such as:

- New services: wavelengths-on-demand, time-of-day service, native-rate local are networks (LAN), 10 GbE, etc.
- Service-level agreements: type of service (ToS), class of service (CoS), and quality of service (QoS).
- Network protection: varying degrees of dynamic wavelength protection and restoration.

Another important capability that can be offered by a wavelength switch is the ability to broadcast the contents of one wavelength to multiple other wavelengths. This functionality is key in applications such as video distribution, and further serves to distinguish a transparent wavelength switch from an O–E–O switch.

FIGURE 1

The Wavelength

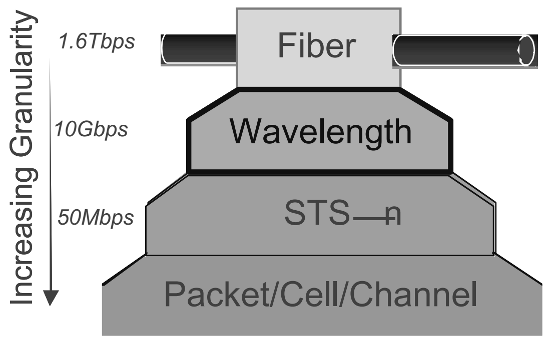

FIGURE 2

The Wavelength Switch

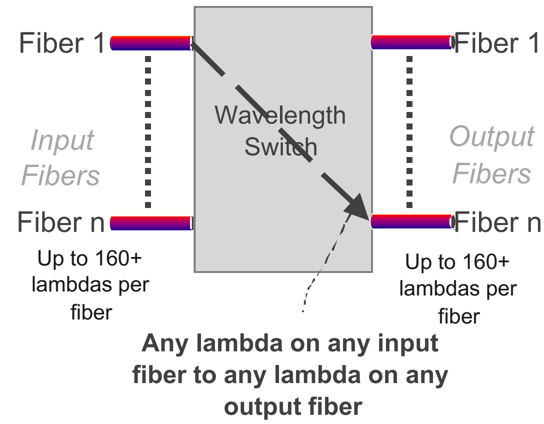

The Evolution of Wavelength Switches

Many of the components required to build a fully functional, all-optical wavelength switch are based on state-of-the-art technology that is in the process of moving from lab to product. For this reason, the first all-optical wavelength switches will not reach the market until the end of 2001.

In the meantime, a number of interim approaches are being proposed. The following sections compare these approaches with the "ideal" switch described in the previous section.

First-Generation Wavelength Switches

In today's telecommunications networks, the basic functionality of a wavelength switch is achieved by combining multiple discrete products (see *Figure 3*). For simplicity, only one direction of signal flow is represented, and passive optical components such as splitters, combiners and filters have been omitted.

Following the signal flow from left to right, each DWDM wavelength on the incoming fiber is terminated, and then, typically, converted to a generic short-haul 1,310 nanometer (nm) interface by an outboard O–E–O transponder. Because each DWDM equipment vendor uses the International Telecommunication Union (ITU) grid in a different manner, the transponders are vendor specific.

The 1,310 nm optical signal is then passed to the optical switch, where it is converted back to electrical, demultipexed into a stream of sub-rate signals, and cross-connected to the appropriate output port using an n x STS–1 matrix. The signals are then multiplexed back to a higher SONET rate and converted back to the 1,310 nm optical signal—again using an O–E–O device. (The label "optical" here is clearly a misnomer, as the core switch is electrical, and for this reason such switches are often referred to as an O–E–O switches.) Finally, the 1,310 nm output from the switch is passed to another outboard O–E–O transponder, which terminates the 1,310 nm signal converts the signal to the appropriate DWDM wavelength.

Today's O–E–O switches utilize an "n x STS–1" based electrical core, and the O–E–O transponders are clocked at the

FIGURE 3

First-Generation Wavelength Switches

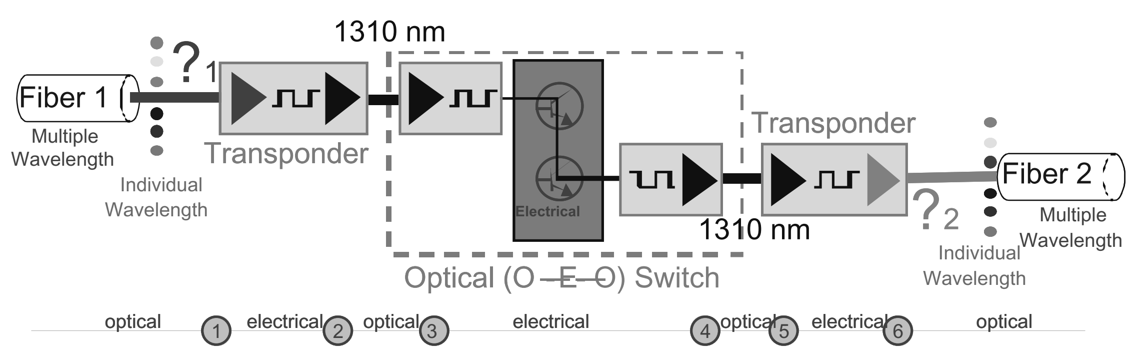

SONET rate; hence neither device provides protocol or bit-rate independence. This architecture is therefore referred to as "opaque" to distinguish it from a transparent wavelength switch.

This first-generation architecture is a very crude attempt at wavelength-switch functionality and suffers from numerous drawbacks:

- Changes in bit rate (i.e., OC–192 to OC–768) require major hardware change outs.
- The architecture is restricted to use with SONET traffic.
- The large number of optical-to-electrical conversions (six per wavelength) makes first-generation wavelength switches expensive, low density, and power hungry.
- Electronic-switch cores do not scale well and require multiple inter-connected matrices to increase capacity. As a result, the cost and footprint of the core increase almost geometrically as the port requirements increase.
- There is one patch cord for each O–E–O transponder, and one transponder for each wavelength. This results in an extremely large number of optical patch cords with the associated increase in cost, reduction in reliability, and administrative headaches.
- Finally, the optical-interface card associated with each wavelength has a different part number and hence one spare transponder card is required per wavelength. Because spares are typically maintained at every office, the result is a huge inventory of transponder cards with the associated cost and administrative headaches, particularly in systems supporting close to 200 wavelengths per fiber.

Second-Generation Wavelength Switches

Proposed alternatives to fully transparent wavelength switches replace the electrical core of today's O–E–O architecture with an optical matrix (see *Figure 4*). Most of these offerings are slated for deployment in the second half of 2001. Based on recent press releases, optical micro-electromechanical systems (MEMS) appear to be the technology of choice for the optical matrix; however, other options are under evaluation, including liquid-crystal and bubble technologies.

Although the core switches proposed for use in these second-generation offerings are transparent, they have no visibility of individual wavelengths carried by the fiber and hence function as an automated fiber patch panel. In order to switch individual wavelengths, transponders are required in conjunction with the optical-switch matrix.

Second-generation wavelength-switch functionality can be achieved in multiple ways. Depending on the particular configuration, the overall architecture will exhibit some combination of the following advantages and disadvantages:

- All-optical switch cores are in the early stages of development and show promise to scale more cost effectively than their electrical counterparts.
- Comparing *Figures 3* and *4*, the second-generation architecture reduces the number of O–E–O conversions per wavelength from six to four. This simplification is likely to result in improved cost and footprint over the first-generation approach; however, since the improvement came from the removal of low-cost intra-office interfaces and not the more expensive DWDM optics, the overall savings are likely to be modest.
- Changes in bit rate (i.e. OC–192 to OC–768) require major hardware change outs.
- Some architectures are restricted to use with SONET traffic.
- The optical-switch matrix cannot support broadcast capability.
- There is one O–E–O transponder, and one interconnection between the transponders and switch for each wavelength in each direction of traffic. This results in an extremely large number of optical patch cords with the associated increase in cost, reduction in reliability, and administrative headaches.
- The optical-interface card associated with each wavelength has a different part number and hence one spare transponder card is required per wavelength. Since spares are typically maintained at each office, the result is a huge inventory of transponder cards with the associated cost and administrative headaches, particularly in systems supporting close to 200 wavelengths per fiber.

FIGURE 4

Second-Generation Wavelength Switch Architecture

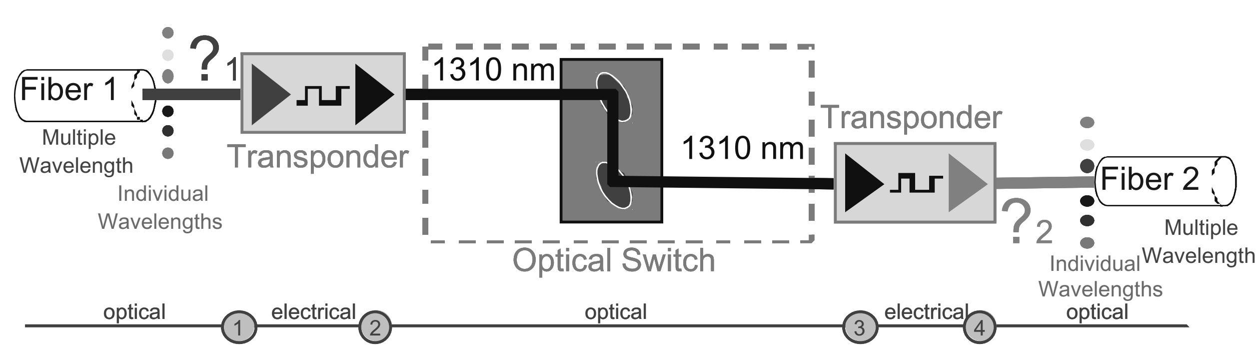

- Direct inter-working with other vendor's ITU grid wavelengths without tunable interfaces is likely to be a major challenge.
- Bundling wavelengths into bands can lead to significant under-utilization of fiber. In some cases over 50 percent of the fiber capacity could be unused.
- Wavelength engineering is a nightmare.
- Switching wavelengths on a band basis is significantly less flexible than switching individual wavelengths.
- The optical switch matrix cannot support broadcast capability.

Third-Generation Wavelength Switches

To meet all the requirements of the "ideal" wavelength switch described earlier in this paper requires a sophisticated combination of tunable optical devices including lasers, wavelength converters, filters, and multiplexers.

The basic architecture of a third-generation all-optical wavelength switch is shown in *Figure 5*. As a wavelength enters from the left a transparent, tunable wavelength translator converts it to a second, fully selectable outbound wavelength. A transparent optical core then directs this new wavelength to the appropriate output fiber thus completing the connection.

In contrast with the first- and second-generation architectures, the third-generation architecture is the first to offer a fully integrated, fully functional solution. In particular:

- The third-generation architecture is 100 percent transparent and hence totally independent of both protocol and data rate.
- Any wavelength on any incoming fiber can be dynamically converted to any wavelength and exit on any other fiber without conversion to electrical.
- All fiber interfaces are at the DWDM level, avoiding the need for the mass of per-wavelength patch cords required to interconnect today's architecture and some second-generation architectures.
- Wavelengths can be switched individually, avoiding the stranded wavelengths associated with the "bundled" approach.
- Tunable optical devices are still early on the technology curve and hence have huge potential for cost reduction and miniaturization.

FIGURE 5

Third-Generation (All-Optical) Wavelength Switch Architecture

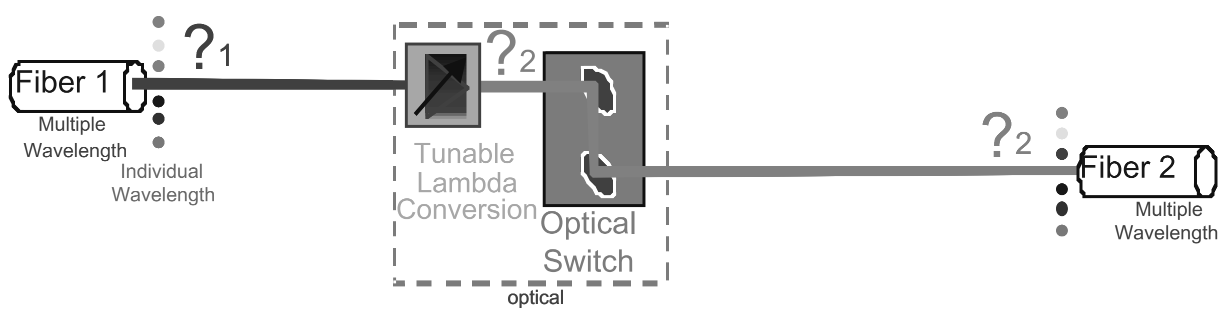

FIGURE 6

A High-Level Comparison of Alternative Wavelength Switch Architectures

	First Generation	Second Generation	Third Generation
Availability	Now	3Q01	4Q01
Protocol Independence?	×	-*	✓
Bit-Rate Independence?	×	-*	✓
Per-Wavelength Switching	✓	-*	✓
Fiber Utilization	✓	-*	✓
Fiber Count	✓	-*	✓
Scalability?	×	-*	✓
Cost/Footprint	×	-	✓
Broadcast?	×	-*	✓
Spare Interface Cards	×	×	✓
Per wavelength Patch Cords?	×	×	✓

× Bad - Average ✓Good * Varies according to switch architecture

- The signal carried by any incoming wavelength can be broadcast to any number of outbound wavelengths.
- Wavelength engineering is simple, even in very large networks.
- The signal carried by any incoming wavelength can be broadcast to any number of outbound wavelengths.

Conclusion

Figure 6 summarizes the relative merits of these architectures. While the second-generation proposals appear to offer some advantages over today's first-generation approach, given the short time until fully functional third-generation wavelength switches become available, these advantages will most likely be short lived and the associated investments stranded.

Perhaps more importantly, in an environment where bandwidth prices are dropping at an astonishing rate (reportedly as high as 50 percent every six months), services providers can no longer compete on price alone. Instead, the ability to introduce new services quickly and effectively will be the key to survival. As described earlier, transparent routing is both bit-rate and protocol independent, providing a sound platform for new-service capability and increased service velocity.

The exciting news is that fully transparent, third-generation wavelength switches are finally viable.

MPLS Applications in Broadband Networks

Anna Wielosz
Associate Director, Broadband Technology, ATM Network Architecture
Bell Canada

Abstract

Switching-based broadband networks represent significant capital investment for many network operators. These networks are used for multiple services and run over multiple protocols. Today, many of them have the possibility of turning on multiprotocol label switching (MPLS) capabilities. This paper analyses the opportunities that the MPLS brings to carriers on the existing networks.

The subject can be divided in two aspects: core and access networks. The core network analysis is focused mainly on the evolution toward a next-generation MPLS–enabled architecture. However, MPLS helps also to improve existing network performances by introducing packet interfaces and label-switched path (LSP)–based management. Access networks are analyzed from provisioning perspective: moving from permanent virtual circuit (PVC) to LSP. An additional aspect for MPLS at the edge is the new service potential that it creates: virtual private networks (VPNs) and voice over MPLS.

Finally, multiple standardization organizations have standard definition works in progress that will allow interoperability between the MPLS and existing Layer-2 technologies. Some of them, such as MPLS user network interface (UNI), may play a crucial role in making MPLS successful at the edge and thus becoming a truly end-to-end protocol.

MPLS Introduction

MPLS has been initially developed to address the following aspects: traffic engineering, quality of service (QoS), and routing performances in Internet protocol (IP) networks. The concept of the MPLS is the following: the IP forwarding based on IP addresses is replaced by label switching. IP routing still takes place, but the label switch routers (LSR) distribute and assign labels, and bind them to IP addresses. Once this assignment is done, no more header look-up is required, but packet forwarding is performed based on label switching. All MPLS packets carry at least one label information.

LSPs are established between label edge routers (LERs) using routing protocols. LSPs introduce a connection-oriented aspect into connectionless IP networks. Moreover, when LSPs are established, resource availability may be verified and the path may be routed according to the QoS requirements. This is a significant improvement of the IP network engineering.

Another important aspect is the underlying protocol independence. MPLS is designed to run over asynchronous transfer mode (ATM), Ethernet, frame relay (FR), and packet over synchronous optical network (SONET), or POS, links. This addresses a crucial challenge in telecom networks: interoperability between technologies.

Today, we see large IP networks deploying successfully MPLS technology and implementing VPN services. For a while now, the MPLS has been considered as a solution for Layer-2 networks. Like ATM a few years ago, MPLS looks promising to be a convergence networks protocol. In addition, today we know that most of the networks carry mainly IP traffic, which eventually should favor the success of the MPLS.

Industry Trends

The first MPLS products on the market were large IP routers conceived for IP backbones. Today, we see more and more ATM and FR switches implementing MPLS protocol as well. MPLS has introduced connection-oriented approach and QoS on IP platforms. The traditional ATM and FR products are by definition connection-oriented. They also have backplanes and switching fabrics supporting hardware implemented QoS. This is a huge advantage for the manufacturers of these products. Well-known QoS mechanisms based on separate queues and exhaustive round-robin algorithms provide very deterministic QoS behavior. A second advantage is the resource reservation algorithms built in traditional switches. The resource reservation constitutes a main tool for traffic engineering and QoS. The starting point toward MPLS for switching platforms is far advanced in comparison to the routers; most of hardware requirements are already implemented and well known.

It should be noted that most of the traditional switches today support IP–over–ATM or FR protocols. Thus, implementing MPLS represents the extension to already existing features. This is effectively the way in which most of the traditional switches evolve today.

Service providers operating large ATM networks and various access technologies realize that today their networks have the capability to implement MPLS. Some of them have already built IP backbones in parallel.

What kind of applications and benefits could MPLS bring to the existing networks? First, MPLS is new services enabler on broadband networks. Layer-3 VPNs can now be created on already-deployed networks and generate new revenue opportunities from already-invested capital. However, the trade-off between service definition on IP–focused equipment and IP–enabled equipment may make this opportunity less appealing. New services based on MPLS standards, such as MPLS Layer-2.5 VPNs or voice over MPLS, will create more attractive service opportunities. Still, the success of the approach depends on the adaptation of MPLS standards by end-user and edge-equipment manufacturers.

There are other appealing aspects for network operators brought by MPLS. A closer look at the access and core networks shows that there are many challenges that network operators are facing today. We will go through some MPLS features and see how they can resolve some of broadband network issues.

Broadband Networks Challenges

Broadband networks (see *Figure 1*) are truly multiservice today. ATM networks carry high-quality video services as well as best-effort Internet services. However, there are still parallel core networks required for Ethernet, Gigabit Ethernet services, and IP services. The reasons behind that? The ATM network did not catch up with required access speeds and is not optimized for best-effort services with Ethernet or IP accesses. Other technologies, such as resilient packet ring (RPR), are adapted better and provide better implementation costs for these types of services.

The second issue is the network scalability. Services such as transparent LAN or Layer-3 services require static PVC mesh. Dynamic protocols, such as LAN emulation, can be applied on small-enterprise networks but are not scalable for large multi-customer networks. The number of PVCs on ATM core grows very rapidly and reaches thousands of connections. Hence, the management of the network becomes a challenge. Any change to a service description may require intervention to multiple PVCs. Similarly, bandwidth allocation on PVC–based networks requires micromanagement. SVCs may address this type of issues. However, because the edge-equipment implementation did not follow this direction, the demand on the SVC type of services has been low.

Bandwidth efficiency is another argument very often used against the ATM network. The 53-byte cells allow deterministic QoS but introduce so-called cell tax, which may reach 30 percent, depending on the packet size distribution. This percentage becomes non-negligible when optical carrier (OC)–48 and higher speeds are concerned.

Additionally, in traditional networks, the interworking functions are required per PVC to adapt FR, Ethernet, IP, etc. This involves extra provisioning efforts for each adaptation circuit.

Some of these issues may be improved by MPLS deployment in the core, while others may require applications in the edge. Because core and access implementation can be done independently, both applications will be described separately in the following sections.

Core Network MPLS Applications

MPLS promises the possibility to collapse multiprotocol core networks. Hence, MPLS–based equipment could be

FIGURE 1

Broadband Core Network Architecture

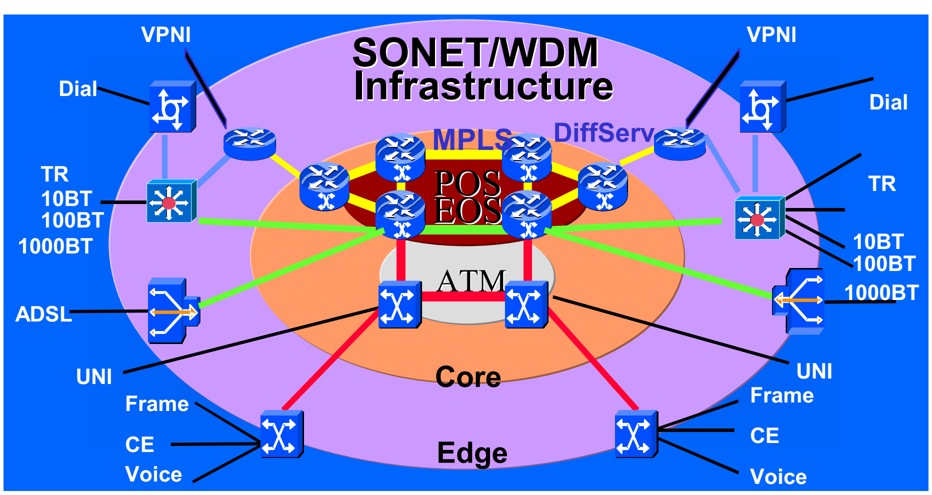

Source: Bell Canada

used as well for IP and for ATM transport. The Internet Engineering Task Force (IETF) and MPLS Forum contributions on Layer-2 mediation and VPNs over MPLS allow adapting the Layer-2 traffic on the Layer-2.5 core. The main advantage is that the convergence is possible without major capital investment in today's networks.

Collapsing multiple networks is driven by investment efficiency requirements. Building parallel backbones result in running two underutilized backbones. This contributes largely to the increase in network costs. Moreover, the return on capital on separate networks is much slower.

Another requirement is the large fan-out to fill network trunks. On average, access links are utilized at 30 percent of the access speed. Because the switches do not allow any backplane oversubscription, the networks need to be, in general, hierarchical in order to provide accurate trunk utilization. On the other hand, the hierarchy also allows the building of less network trunks and thus better fiber utilization (see *Figure 2*).

The hierarchical network architecture gives an opportunity to evolve IP and ATM backbones to a common core architecture. This evolution can be done without significant changes to the existing networks. The new common core should provide high-speed "pipes" with the support of LSPs. Current backbones will perform mediation function and will aggregate traffic into LSPs.

The common MPLS core will have very stringent requirements for QoS functions. ATM backbones should be able to receive the same grade of service satisfying real-time application requirements. With the switching fabrics and queuing mechanisms implemented in next-generation MPLS switches, it becomes a reality.

The constraint-based routing–label distribution protocol (CR–LDP) and resource reservation protocol–traffic engineering (RSVP–TE) protocol, both label-distribution–related MPLS mechanisms, provide the implementation of resource reservation and QoS mechanisms. This way, LSPs are using routes where QoS requirements can be satisfied.

In the MPLS core, traffic-engineering functions will be pushed to the edge. The traffic over an LSP will be admitted according to its QoS and bandwidth requirements (forward error correction [FEC] and connection admission control [CAC] algorithms). The edge will be also responsible to provide appropriate traffic aggregation on an LSP. Thus, all flows will be policed at the edge: IP, ATM, or frame-based rate limiting. Secondly, the number of flows expected over an LSP should be engineered in the access. This approach gives a possibility of intelligent overbooking based on the community of interest.

Once this classification is done at the edge, the core is managed on a "big pipes" basis. This approach should improve network resources provisioning. An LSP bandwidth reflects well the requirements of the community of interests it interconnects. LSPs are then modified as requirements grow or, ideally, adjust bandwidth to traffic in a dynamic fashion (see *Figure 3*).

MPLS promises also performing fast rerouting procedures. The rerouting times are expected to be in the range of 100 ms. This significantly improves the robustness of core networks. Today, Layer-1 robustness does provide link-failure or card-failure but not node-failure recovery. Layer-2 and Layer-3 rerouting mechanisms' performances are in the range of seconds.

With broadband core network growth in the range of 20 percent and 30 percent for the next two years, more links are expected to be at OC–48 and higher speeds (see *Figure 4*). For some time now, the OC–3 component of the total core bandwidth has decreased steadily, leaving a place for OC–12 and higher speeds. Similar behavior will be

FIGURE 2

Hierarchical Network Architecture with MPLS Core

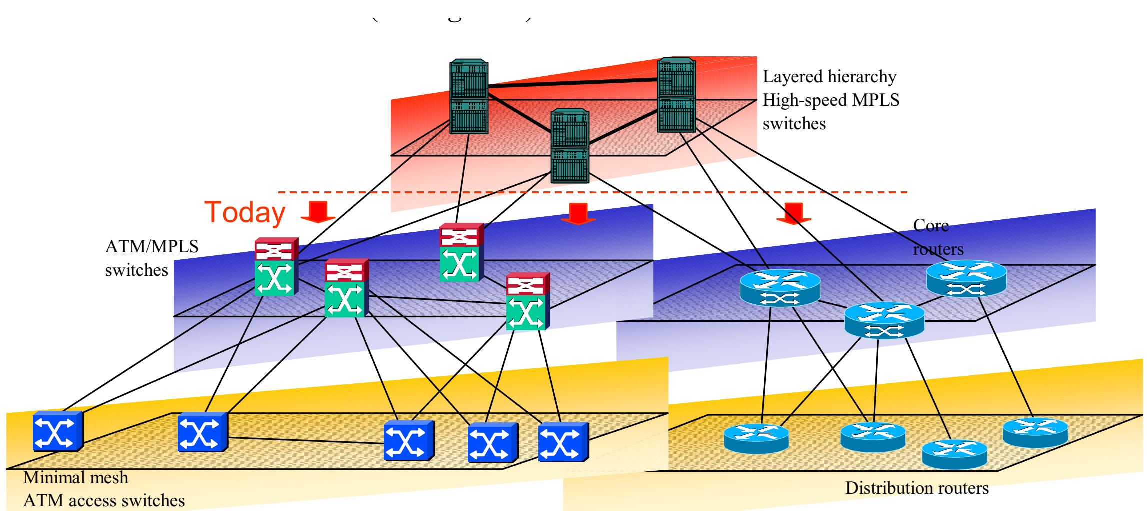

FIGURE 3

LSP–Based Bandwidth Management

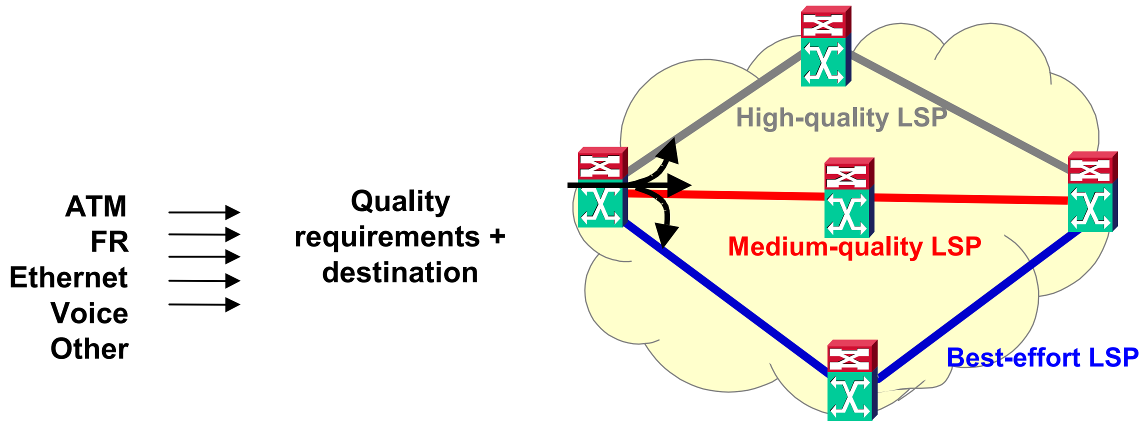

FIGURE 4

Link-Type Percentage in Total ATM Core Bandwidth

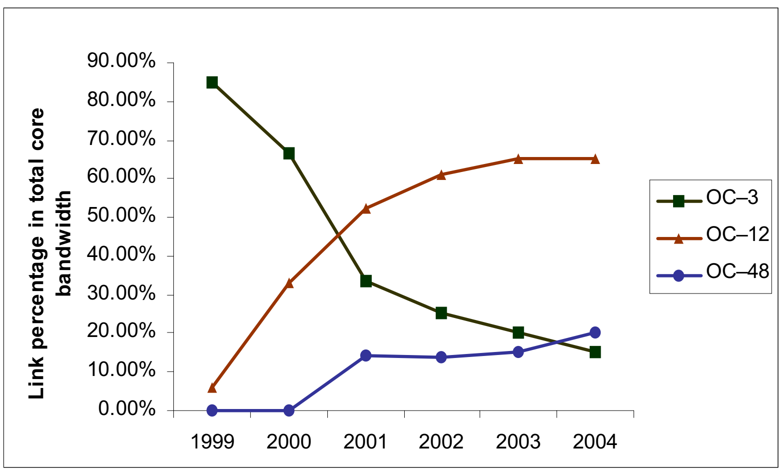

Source: Bell Canada

observed with OC–12 links. The number of higher-speed links is expected to grow. This is the result of growth applied to a total network size and growing bandwidth demand per circuit.

First, it is beneficial to deploy higher capacity as early as possible. Otherwise, the deployment of parallel links grows with capacity requirements. Moreover, redesigning the network with the proper trunk speeds involves significant operational effort. Once again, collapsing the network helps in deploying the right infrastructure at the beginning.

Finally, it is worth to look at the bandwidth savings with packet technologies, especially for high-speed interfaces. Jitter introduced at these speeds is lower, and maximum acceptable packet length becomes higher. A very interesting feature under development by the MPLS Forum is the cell concatenation. Multiple cells are concatenated into an MPLS frame without analyzing any payload information. This allows the reduction of cell tax without performing the complete reassemble at the packet level (see *Figure 5*).

Replacing cell interfaces by packet or frame interfaces, assuming that the headers are implemented as per *Figure 5*, can produce up to 16 percent of bandwidth savings. This is equivalent to 398 megabits per second (Mbps) per OC–48 link. These interfaces can be POS, Gigabit Ethernet, FAST, or other interfaces supporting MPLS.

In summary, implementing MPLS in broadband core networks gives a clear evolution path toward a common core next-generation network. Management on large LSP bases can provide significant improvements in bandwidth management. Finally, the smooth introduction of packet interfaces with QoS support is enabled, which translates directly into significant bandwidth savings.

FIGURE 5

Proposed MPLS Header for ATM Mediation as per....

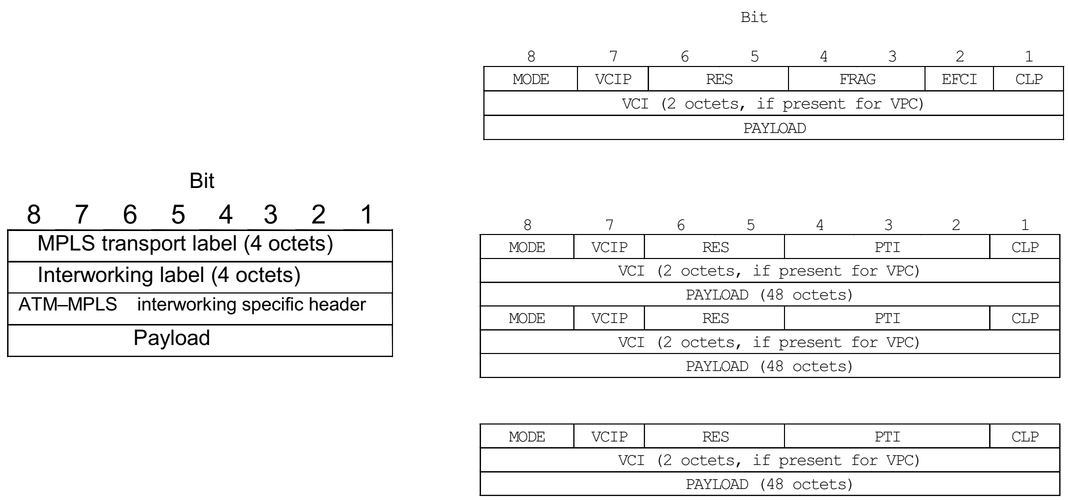

Edge MPLS Applications

Independently to the core, edge platforms can be deployed with MPLS capabilities. Assuming the broadband core can switch or tunnel the LSPs, the following advantages should be considered: provisioning advantages coming from circuits' aggregation and new services opportunity.

As has been mentioned, multiple circuits in fully meshed configurations must be provisioned today to provide services on ATM networks. MPLS does not resolve the problem completely; routers still need to be interconnected with LSPs. However, LSPs are established automatically according to IP routing protocols, and LSRs are exchanging topology information.

Furthermore, even if the mesh is still required, the degree of meshing is significantly reduced by the aggregation step. Because IP headers are processed at the access switch acting as the LER, multiple access links can be mapped on the same LSP. It reduces significantly operational impacts and time to market for new services. Enabling a new access does not require any intervention in the core network (see *Figure 6*).

The interoperability problem is reduced from any-to-any to IP over MPLS, where IP can be carried over any Layer-2 protocol.

The common IP control plane plays a role in automatically establishing LSPs across a heterogeneous network. Like S–PVC in ATM networks, LSPs are established via signaling (LDP). The topology information is exchanged, and optimal path in the network is selected according to the LSP parameters. Both, CR–LDP and RSVP protocols contain LSP parameters: bandwidth and QoS requirements. For more information on QoS, please refer to [1].

Additionally, MPLS edge implementations will become new service enablers. Today, the MPLS forum is working on a voice-over–MPLS standard. Considering the MPLS QoS capabilities, that could become more successful than voice over IP (VoIP) applications. Another interesting standard under development is the MPLS UNI. The standard is meant to provide VPN capabilities with autoconfiguring edge.

MPLS Standards Development

The standards that are crucial for the success of MPLS in the converged backbones are those defining mediation policies—i.e. ATM, frame over Layer-2.5 networks. Major equipment manufacturers are involved in developing these specifications, and first implementations are available on the market.

Mediation is the technology allowing smooth migration or integration of existing Layer-2 networks with MPLS networks (see *Figure 7*). MPLS networks also promise interoperability with a generalized MPLS (GMPLS) optical core. The ATM Forum, FR Forum, and MPLS Forum work is focused on the transport of ATM or FR over MPLS. At this time, any-to-any Layer-2 interoperability is not considered. However, by definition it is possible with IP routing.

Layer-3 VPNs have been addressed by IETF Internet drafts [2]. Layer-2 transport is first addressed by "Martini" drafts describing Layer-2 MPLS encapsulation—i.e., ATM, FR, Ethernet, point-to-point protocol (PPP), and VLAN over MPLS. These drafts specify what information should be included in the headers and what management information should be exchanged (such as P–NNI, FR link management, etc.). A more complete Layer-2 VPN solution is proposed in the "Kompella" draft and extensions to it. These contributions propose extensions to LDP protocols to carry CE identification information [3] [4].

Another contribution under development is the UNI MPLS. This work may provide an attractive way to provide MPLS VPNs with an autoconfiguring feature for customer-premises equipment (CPE) devices. The CPE discovers LSPs on

FIGURE 6

MPLS Edge Implementation

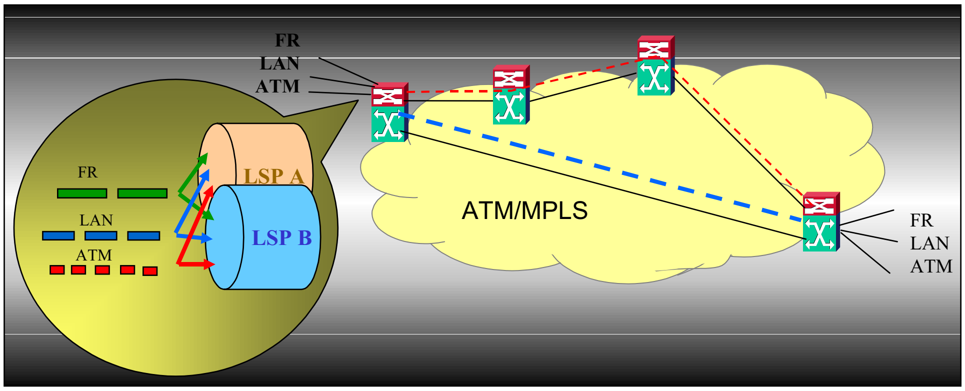

FIGURE 7

MPLS Mediation

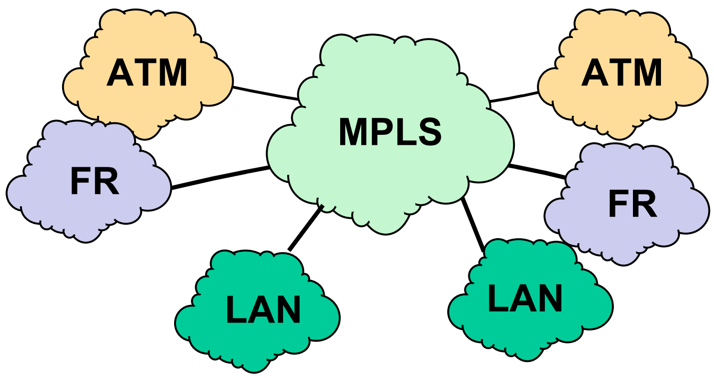

the link to which it is connected and configures its forwarding tables accordingly. VPN information is exchanged between the LER and CPE device over LDP.

The important limitation to note is that LSPs are presently point-to-point and unidirectional connections.

Conclusions

In conclusion, broadband networks can evolve toward convergence with IP technologies by implementing MPLS in the core. MPLS implementation will generate a profit over the already invested capital, as the functionality is incorporated on most of traditional switching platforms.

Multiple advantages appear in the hierarchical structure of a network. It enables almost a seamless implementation of next-generation MPLS cores shared between traditional switching and IP backbones. This, in turn, allows better resource utilization and, therefore, more efficient investment.

Additional advantages are related to operating core networks with a low number of high-speed trunks. The n-square problem of ATM networks, although still present, is reduced to the number of devices necessary to interconnect. The aggregation over an LSP reduces the complexity of new services provisioning.

Finally, MPLS brings new service opportunities to the edge. It also enables a huge installed base of broadband switching networks to carry IP services.

References

Please refer to the ATM Forum, MPLS Forum and IETF Web sites for specific contributions: http://www.atmforum.com/, http://www.mplsforum.org/, and http://www.ietf.org/

1. Paul Brittain and Adrian Farrel, "MPLS Traffic Engineering: A Choice of Signaling Protocols," *Data Connection*, January 2000.
2. E. Rosen and Y. Rekhter, "BGP/MPLS VPNs," RFC2547, March 1999.
3. "Encapsulation Methods for Transport of Layer 2 Frames over IP and MPLS Networks," Internet Draft, July 2001.
4. "MPLS Based Layer 2 VPNs," Internet draft, June 2001.

Subcarrier Multiplexing: More Than Just Capacity

Ted Wolcott
Director, Advanced Optical Product Development
Centerpoint Broadband Technologies

Heidi Thelander
Senior Product Marketing Manager
Centerpoint Broadband Technologies

Abstract

Increasing demand for bandwidth in today's networks has driven a rash of optical technology development. Among these new technologies are high-capacity subcarrier multiplexed systems capable of increasing network capacity as much as eight times while dramatically improving fiber bandwidth efficiency. However, this subcarrier multiplexing (SCM) technology does more for the network operators than just provide greater capacity and higher efficiency. This paper describes how the use of higher order modulations together with SCM provide inherent capabilities and features which are highly attractive to network operators when implemented effectively.

1. Introduction to Subcarrier Multiplexing

Demand for more flexible bandwidth creation at lower cost is causing service providers to look at alternative and complementary technologies to traditional time division multiplexing (TDM) and wavelength division multiplexing (WDM) systems. SCM provides an additional dimension of multiplexing to increase the efficiency and flexibility of optical transport networks. It has largely been promoted as an approach for increasing capacity and bandwidth efficiency on a given fiber or wavelength with its use of bandwidth efficient digital modulation and sublambda grooming. However, SCM and the architectures it enables provide many additional capabilities for which network planners are waiting.

SCM, sometimes referred to as frequency division multiplexing (FDM), has been used in radio, satellite, and cable television (CATV) applications at much lower data rates for many years, providing low-cost, highly bandwidth-efficient physical-layer transport. Recent developments in high-rate digital signal processing and radio frequency (RF) component capabilities have enabled application of this technology to optical networking systems, which have just begun to require such high levels of bandwidth efficiency. This optical SCM technology is the result of the combination of several key elements, all considered proven technology in their other applications:

- Higher order modulation for bandwidth efficiency
- Pulse shaping for spectral efficiency
- Forward error correction and coding for bit-error rate (BER) performance
- RF signal multiplexing and demultiplexing
- Broadband optical intensity modulation

The higher order modulation most commonly used in SCM is quadrature amplitude modulation (QAM) wherein multiple information bits are transmitted in a single signaling interval using a mapping algorithm involving both amplitude and phase. An example of this is 16-QAM (see *Figure 1a*), which transmits 4 information bits by transmitting one of 16 locations in the amplitude phase plane or "symbols." The digital clock speed used in timing these symbols is much lower than the pure bit clock speed used in an on-off-keyed (OOK) modulation, translating to lower speed (and cost) electronics and to an ability to transmit more information in a given bandwidth. Thus, these modulation types are often referred to as bandwidth-efficient modulation (BEM).

In addition to transmitting several information bits in each symbol period, digital signal processing is used to shape the spectrum of the transmitted pulses so that most of the modulated signal energy is retained in a minimum frequency band. This reduces required spacing between carriers significantly (see *Figure 1b*), enabling subcarrier multiplexed systems to provide very high bandwidth efficiency. Since having 16 states (symbols) to choose from each period would naturally increase the likelihood of error, techniques to enhance performance are also used with these modulation approaches so that link budgets comparable to or better than alternative approaches can be achieved with similar power. These include linear equalization of the received signal within the demodulator and use of multiple advanced forward-error-correction coding algorithms (see *Figure 1c*).

SCM technology is more than digital processing; the key to many of the performance benefits is actually the RF processing that is performed on the already modulated signals. These individual channels—all output from digital modems at a single intermediate frequency—are frequency converted and filtered to a series of equally spaced offset sub-

FIGURE 1

Bandwidth Efficient Modulation Forms the Basis for Optical SCM

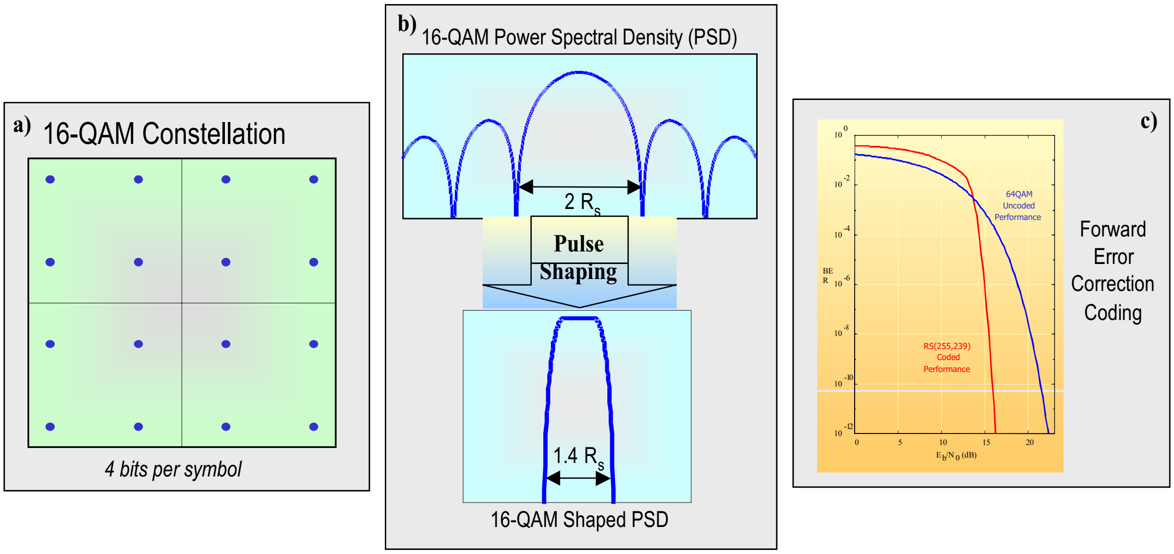

a) 16-QAM Transmits 4 Information Bits Each Signaling Period
b) Spectral Shaping Narrows Bandwidth Required for Each QAM Signal
c) Forward-Error-Correction Coding Improves System BER Performance

carrier frequencies (f_1, f_2, … f_k). The offset subcarriers are combined or multiplexed together, creating a single composite RF signal consisting of *k* independent subcarriers (see *Figure 2a*). This RF signal is then used to intensity modulate the laser, creating a single wavelength optical signal carrying *k* independent subcarrier channels (see *Figure 2b*).

Systems have been developed with these technologies that provide up to 20 gigabits per second (Gbps) of data transmission on a single wavelength in as little as 32 gigahertz (GHz) of spectrum (see *Figure 3*)—equivalent synchronous optical network (SONET) TDM systems require 60 GHz. With linear external modulators, standard International

FIGURE 2

SCM Aggregates Multiple BEM Signals onto a Single Optical Carrier

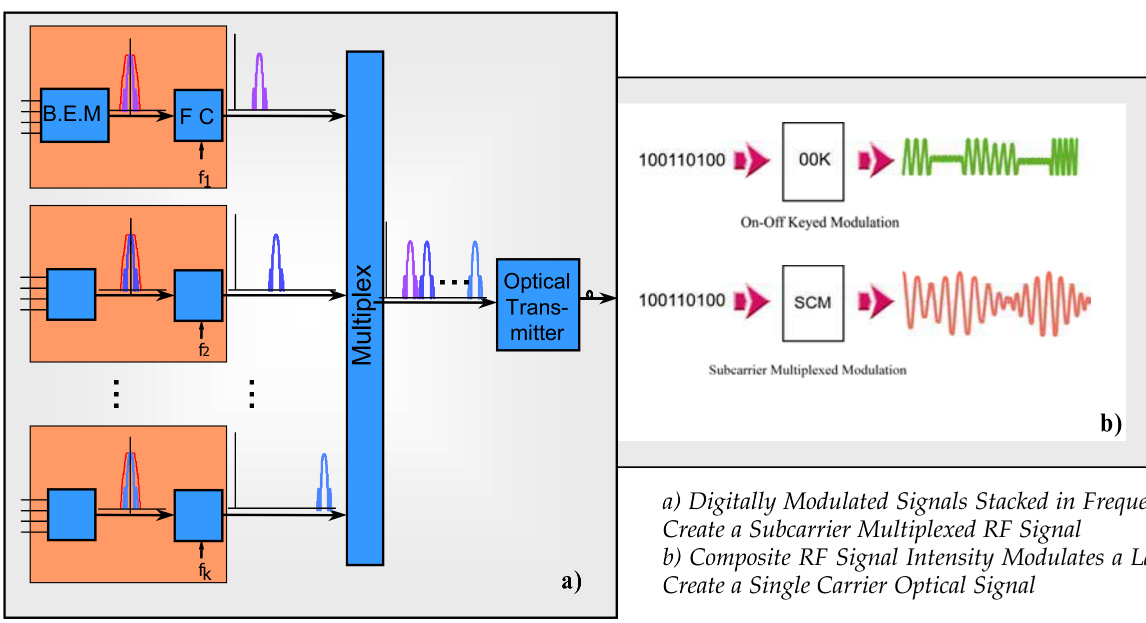

a) Digitally Modulated Signals Stacked in Frequency to Create a Subcarrier Multiplexed RF Signal
b) Composite RF Signal Intensity Modulates a Laser to Create a Single Carrier Optical Signal

Telecommunication Union (ITU) grid lasers are used to transmit these signals at any desired ITU wavelength. Systems have been constructed providing 32 subcarriers capable of transmitting 32 optical carrier (OC)–12/OC–12c signals or eight OC–48/OC–48c signals on a single wavelength. In addition, these individual wavelength signals can be fed into WDM systems to provide even greater capacity per fiber.

SCM provides superior bandwidth efficiency and overall capacity to the much simpler OOK systems. However, capacity and bandwidth efficiency alone will not answer many network operators' needs. Many other functions and features that operators require, desire, or hope for are provided by the unique combination of the digital processing BEM and SCM. Some inherent features that provide a new level of flexibility are as follows:

- Linear and ring configurations enable add/drop, drop/continue, regeneration, and broadcast/drop (fanout) of individual subcarriers
- Transparency providing multiple service types in native format—with carrier-class protection
- Highly flexible traffic mix carries any combination of IP, ATM, and SONET signals
- SONET–like protection performance without SONET overhead on each individual channel
- Interoperability with TDM and WDM equipment, creating flexible, scalable, and robust networks

In addition to the features produced by the architecture, a number of performance benefits are obtained through implementation of this type of system. These include the following:

- Straightforward network link design—no optical link budget adjustments with subcarrier provisioning
- Improved system robustness with reduced adjacent channel interference and link performance enhancements due to equalization and FEC
- Lower susceptibility to polarization mode dispersion (PMD) and chromatic dispersion than 10 Gbps TDM systems
- Highly insightful performance monitoring with many valuable metrics from digital signal processing
- Excellent fault isolation capability due to metrics and architecture

2. Functionality Creates Flexibility

The functionality of an optical networking system determines the level of flexibility it provides to both the network designer and the network operator. Subcarrier multiplexed systems can be architected for a high level of functionality in a number of areas. With low-cost RF processing, the SCM system can be configured to support both linear and ring network configurations (see *Figure 4*). Nodes can be configured as linear terminals, ring add/drop multiplexers (ADMs), or regeneration elements, allowing traffic to be set up as add/drop, drop and continue, regenerated, or broadcast and drop (fanout) for individual subcarriers—and be reconfigured as needed. Linear configurations are then easily created to provide 1+1 protection for application to regional point-to-point interconnect, capacity expansion of thin fiber routes, or to supplement capacity on existing long-haul routes where fiber type currently limits the ability to increase TDM speeds. Metro rings can be configured with ADMs (and regenerators as needed) to provide unidirectional path-switched ring (UPSR) protection for the metro core or collector ring applications where capacity growth is rapid, yet service mix difficult to predict.

In addition to supporting a variety of applications with these configurations, the digital and RF processing inherent in SCM provide a level of visibility together with data transparency that is unique. The digital processing performed yields a great deal of insight into the performance of the link on a given channel with metrics ranging from uncoded channel BER to received signal-to-noise ratio (SNR). This enables SONET–level protection to be provided on *any* traffic type carried by an SCM system. Since all traffic is digitally modulated, all traffic can be protected with less than 50 ms switching to a protect channel, including native IP data traffic without SONET framing or digital wrappers.

FIGURE 3

SCM Packs More Information onto an Optical Carrier

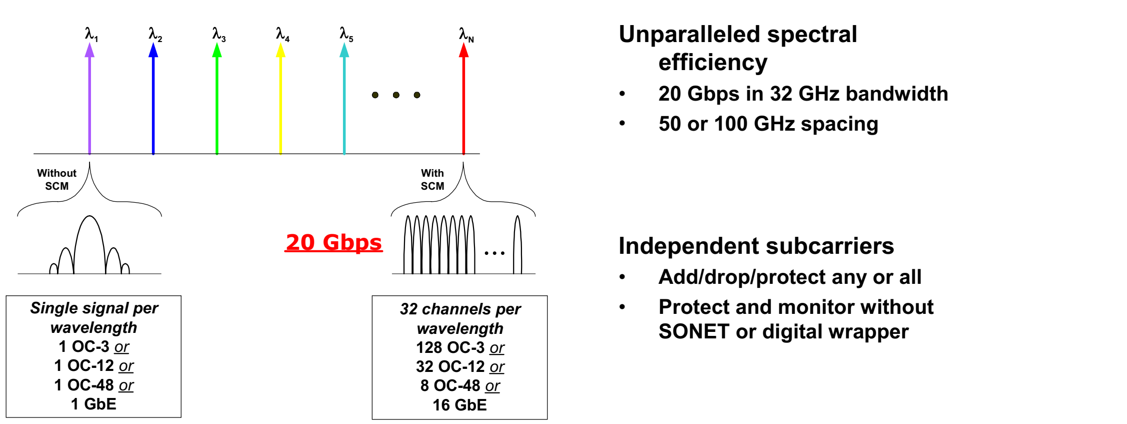

FIGURE 4

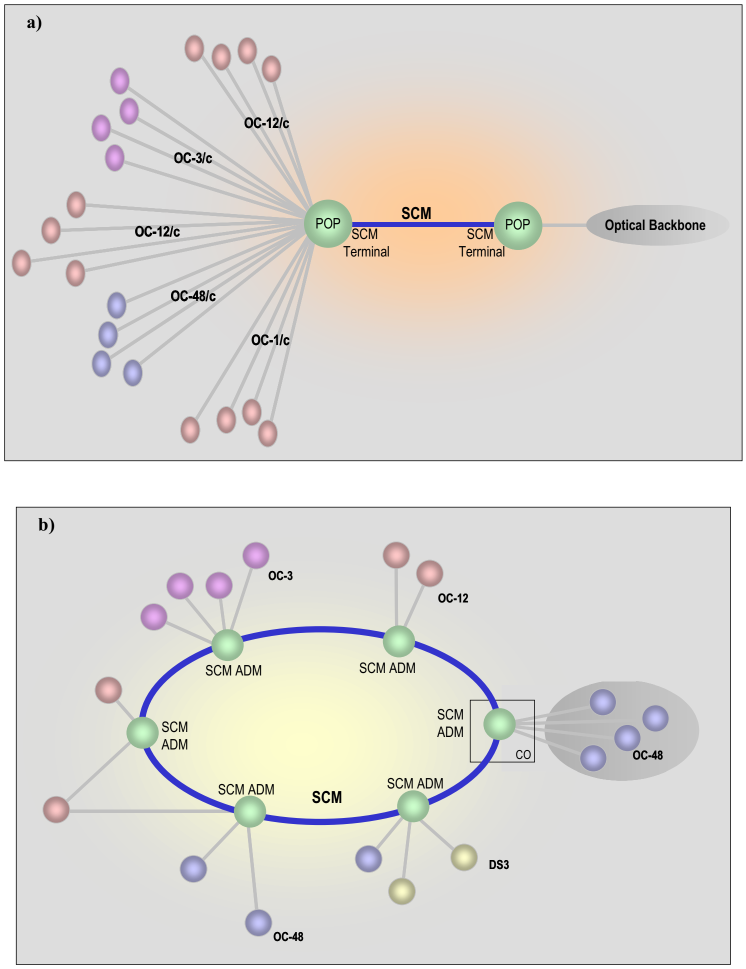

The RF processing performed in the SCM system provides the true transparency sought after by network operators. Each channel can independently carry different traffic operating at different rates (see *Figure 5*). Any traffic type can be carried transparently, and systems have been designed that carry SONET, ATM, and IP traffic seamlessly on the different subcarriers. This transparency provides the network operator with the ability to reconfigure the system easily and quickly as the service mix changes without swapping out a lot of equipment. The rise of Gigabit Ethernet as a service provided by carriers is pushing demand for flexible systems that enable transitions, and SCM eliminates the need for equipment to convert non–SONET signals into SONET signals.

Because SCM is effectively orthogonal to TDM and WDM, the technologies are inherently interoperable (see *Figure 6*). SCM adds another dimension to the space network operators can choose from to implement capacity. On the tributary side, SCM systems can, and indeed must, be designed to provide complete compatibility to SONET TDM signals. The use of standard lasers supporting the established ITU grid wavelengths enables SCM systems to be combined with WDM for very-high–capacity systems with superior flexibility, providing subwavelength granularity while supporting multiple services in their native format.

3. Performance Provides the Edge

With so many technologies being applied to optical networking, traditional notions of performance standards are being challenged and adapted. Network design is more challenging at higher rates, over more wavelengths, for longer distances. Any edge in simplifying network design through enhanced or more robust performance can make a real difference to designers—and to cost. SCM with digital signal processing for bandwidth-efficient modulation provides some significant performance benefits to optical networks.

Optical link budgets can be difficult to maintain through a network that, despite standards used for initial design, has a range of channel conditions across its spans; each change can mean having to take a new look at the link budgets. However, the optical power for an SCM signal remains constant, no matter how many subcarrier channels are in service. By modulating a common light source with the aggregate signal, the optical power is a function of the laser, not the number of channels in the signal. Any power not used

FIGURE 5

SCM Aggregates and Transports Full Mix of SONET, ATM, and IP Traffic Transparently

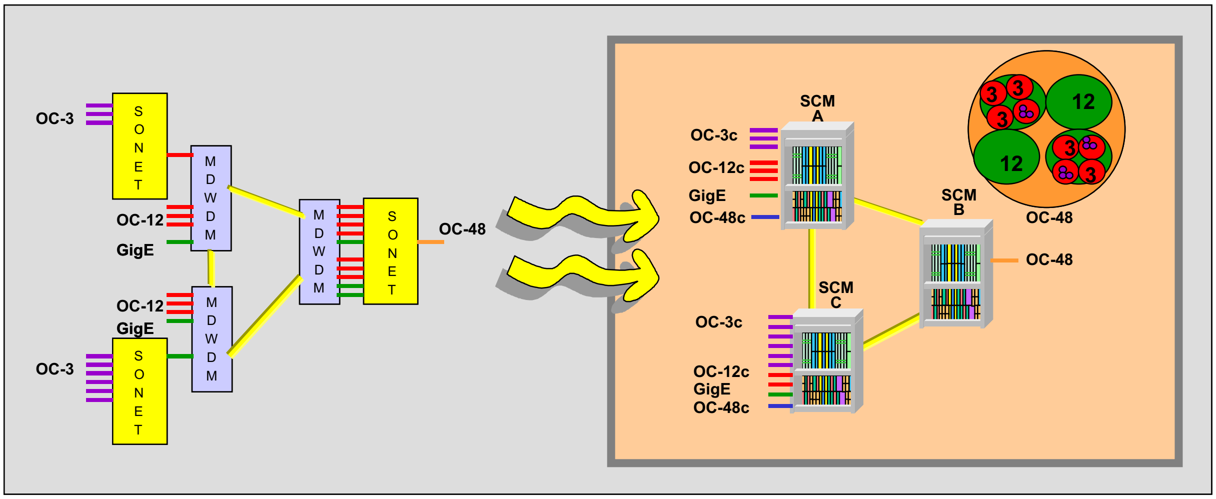

FIGURE 6

SCM Interoperates with TDM and WDM for Three Dimensions in Multiplexing

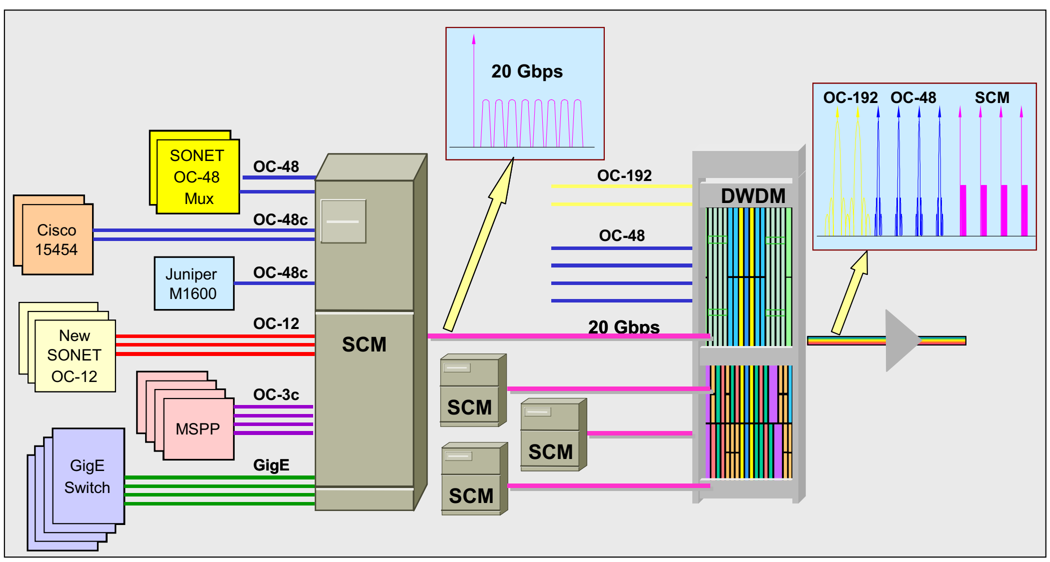

by a channel is returned to the unmodulated portion of the SCM waveform. This enables add/drop of subcarriers without re-engineering the link.

The SCM and BEM signal processing combination also provides a high level of robustness in system link performance over a wide range of channel conditions. Because of the pulse shaping and filtering of the subcarriers, SCM signals have very low adjacent channel interference (see *Figure 7a*). Unlike an OOK signal, the spectral sidelobes of an SCM signal are negligible—an ideal SCM signal will have no sidelobes at all, whereas nonlinearities in any practical SCM modulator will produce small sidelobes well below those of an OOK signal. This lack of out-of-band power results in very low adjacent channel interference in dense WDM (DWDM) systems.

The digital signal processing in the SCM modems provides additional robustness with linear equalization to improve signals received through varied channels and with forward error correction more powerful than that used in other optical systems. The relatively low subcarrier clock rate—a 20 Gbps SCM signal made up of synchronous transport signal (STS)–12 rate subcarriers, operating with parallel channels, each running at a clock rate on the order of 200 megahertz (MHz), rather than 20 GHz—allows the implementation of multi-tap, fractionally spaced linear equalizers at each receiver to combat linear distortion in the channel (see

Figure 7b). The improved signal quality results in lower BER in highly distorted channels.

Many optical systems are now employing linear block code based forward error correction, typically Reed Solomon block codes. However, with binary signaling, the full power of forward error correction cannot be tapped. SCM's use of higher-order modulation allows the traditional Reed Solomon forward error correction to be paired with trellis coded modulation (TCM), a form of forward error correction that takes advantage of the modulation format to maximize the coding gains (see *Figure 7c*). Together, these codes provide coding gains of 7 decibels (dB) or more.

A critical performance edge is found in the area of dispersion tolerance. Chromatic dispersion impairs performance of all high-rate OOK signals, and therefore must be compensated for every link and channel in many optical networks. The higher the bit rate, the wider the signal bandwidth, the lower the chromatic dispersion tolerance. Likewise, the longer the fiber span, the greater the chromatic dispersion. To avoid expensive or complex chromatic dispersion compensation, networks must be limited in either data rate or distance.

Single-sideband (SSB) SCM eliminates the need for chromatic dispersion compensation in all but the longest optical networks, even at very high aggregate bit rates. By eliminating the energy on one side of the optical carrier, the chromatic dispersion tolerance is limited only by the symbol rate of the individual subcarriers—not by the aggregate signal bandwidth. Thus, a 20 Gbps SCM signal, operating with OC–48 rate subcarriers, will exhibit chromatic dispersion tolerance far exceeding that of a 10 Gbps OC–192 system (see *Figure 8*).

PMD is another impairment that gets worse with increased bandwidth. Unfortunately, SSB does not help in this case. However, the inherent bandwidth efficiency of an SCM signal limits the impact of PMD. A 20 Gbps SCM signal exhibits greater PMD tolerance than an OC–192 signal. At 10 Gbps, the PMD tolerance of SCM approaches that of an OC–48 signal. In networks where PMD is the primary limitation, an SCM signal can be designed to maximize the PMD tolerance.

Extensive performance metrics provided by the digital signal processing together with a decentralized architecture provide SCM systems with a very high level of performance-monitoring and fault-isolation capabilities on traffic of all types. Use of TCM and Reed Solomon coding allows the monitoring of both channel BER performance (before decoding) and decoded BER performance (see *Figure 9*). The channel BER is much higher in the operating region providing BER degradation information to the operator as gradual or slight degradations occur, well before they accumulate to result in an outage. This provides a valuable maintenance corrective action tool to avoid costly downtime on service-level agreements (SLAs). SCM systems also provide superior fault isolation and correction as compared to other technologies. Together with proper fault correlation, it isolates problems down to the card level as well as corrects and increases the reliability of the various traffic protocols.

FIGURE 7

SCM and Signal Processing Increase Robustness of Optical Links

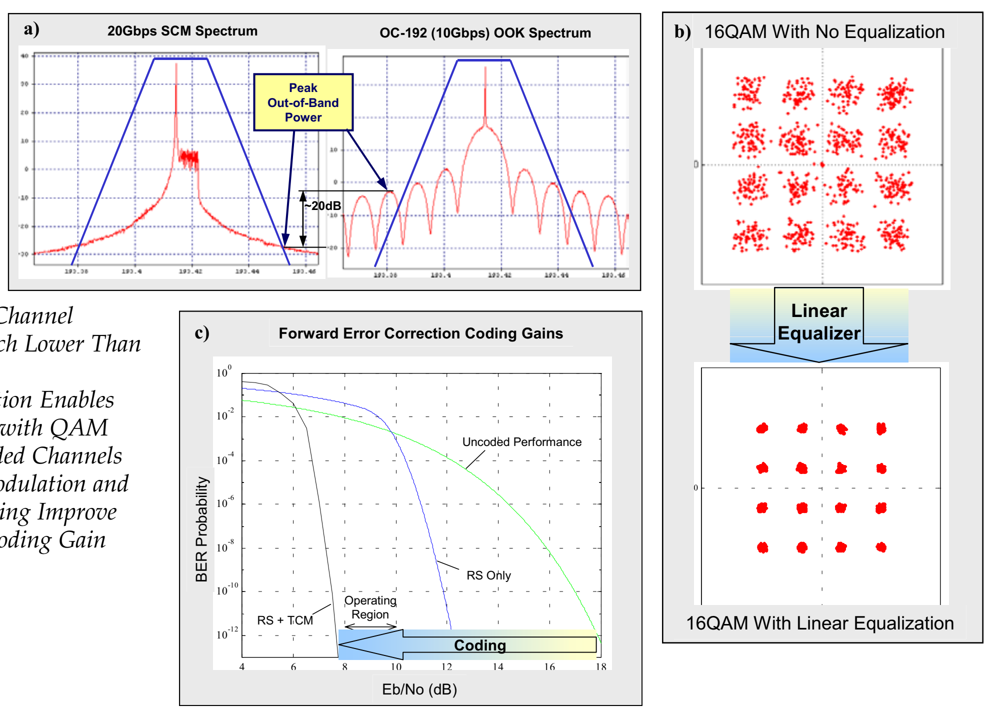

a) SCM Adjacent Channel Interference Is Much Lower Than OOK Systems
b) Linear Equalization Enables High Performance with QAM Signal over Degraded Channels
c) Trellis Coded Modulation and Reed-Solomon Coding Improve Links with High Coding Gain

FIGURE 8

Narrower SCM Channel Bandwidths Dramatically Improve Chromatic Dispersion Tolerance

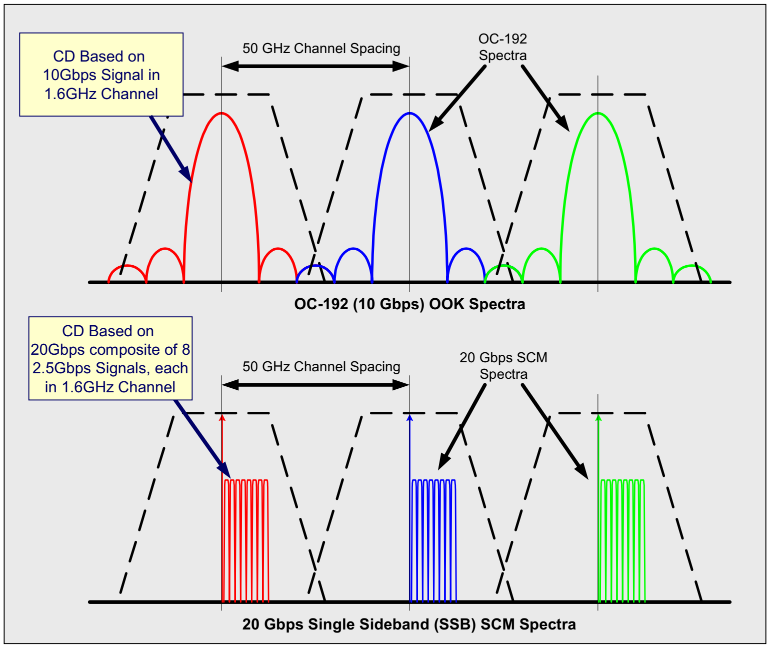

4. Conclusion

While SCM clearly provides the capability to increase optical network capacity and improve fiber bandwidth efficiency, it also provides many other potential benefits to the optical-network designer and operator. The underlying digital signal processing, RF processing, and optical modulation technologies combine to provide a wide range of advantages in the network. Flexibility in configuration and traffic mix combined with true transparency and carrier-class protection of all traffic reduces operations cost and improves provisioning. Highly robust optical links with constant optical power together with highly informative performance metrics and lower susceptibility to dispersions make SCM systems easier and less costly to design and maintain. SCM is clearly part of the future of optical networks—providing more than just capacity.

FIGURE 9

Channel BER Provides Early Detection of Channel Degradations

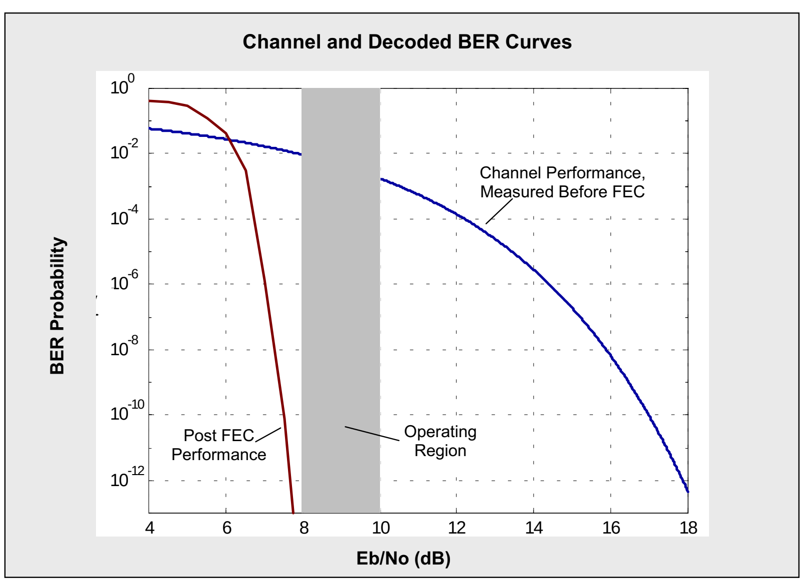

MPLS: Making the Most of Ethernet in the Metro

Tim Wu
Manager, Corporate Development
Riverstone Networks

Introduction

Two of the strongest trends in networking today are the migration of Ethernet from local area networks (LAN) to metro area networks (MAN) and the enhancement of Internet protocol (IP) networks with multiprotocol label switching (MPLS) technology. Most service providers are familiar with either or both of these developments. But they may not be aware of the potential stemming from the combination of the two, particularly for offering premium service-level agreements (SLA).

For service providers building out Ethernet-based networks, or thinking of doing so, the available return on investments is a central question. The kinds of SLAs that a service provider can offer are crucial to these returns. First, SLA capabilities define the scope of potential customers, as certain customers simply will not purchase services below certain standards. Second, the difference between profit and loss on individual deployments may depend on the ability to charge more for services, like transparent LANs, backed by premium guarantees of network performance. For Ethernet carriers, MPLS is the key to the connection-oriented capabilities that translate into premium service-level capabilities. MPLS will, in the long run, provide Ethernet service providers a chance to reach those customers once beyond their ken and improve overall profitability.

This study introduces MPLS and Ethernet technologies and demonstrates in concrete terms how the two can improve the economics of Ethernet-based services using the example of transparent LAN services (TLS) in a typical metro setting.

An Introduction to MPLS and Ethernet

Interestingly, MPLS and Ethernet technologies come from opposite ends of the network (see *Figure 1*). Ethernet comes, of course, from the LAN, while MPLS has its origins in the IP core. Metro networks mark their first and only point of direct interaction.

Ethernet is by far the most highly deployed network type on earth, with over 250 million devices already connected. With its domination of local-area networking secure, Ethernet has most recently taken a well-publicized leap into metro networking. Metropolitan service providers (MSP) take the same simplicity and huge bandwidth capabilities (10-gigabit products are now available; 100-gigabit products are expected within two years) and add long-distance capabilities (40-70 kilometers) and the high-performance, switch-routing equipment suited to carrier environments. The result is a potent combination of speed, scalability, and operational simplicity that has won many converts.

Leading the deployment of Ethernet metro networks are various forms of MSPs, including the Ethernet local exchange carriers (EtherLEC), building local exchange carriers (BLEC), and companies like Telseon, Intellispace, Yipes, and others. But established carriers are also looking to metro Ethernet, particularly as a data-focused supplement to existing synchronous optical network (SONET) and time division multiplexing (TDM) networks.

MPLS has its origins in the opposite end of networking—the network core—where its first deployments (by Global Crossing and UUNet) are found. MPLS has so far been deployed to address various backbone-network problems, such as integrating IP and asynchronous transfer mode (ATM) networks, reducing IP router overhead, and solving route-propagation problems. Yet, because MPLS generally adds connection-oriented, path-switching capabilities to IP networks, it is an extremely versatile technology. In the metro, its most obvious utility comes as an instrument for creating guaranteed and secure-service capabilities.

How the Two Work Together

In many ways, MPLS and Ethernet are an ideal match. Ethernet is well known for its simplicity and comfort level with IP data traffic. Ethernet is also known, however, to lack sophisticated quality of service (QoS) capabilities; this may be the key to its low cost but is also a major shortcoming.

MPLS improves on Ethernet not by changing what is good about Ethernet but by adding desired capabilities at a higher layer (what some call Layer 2). When MPLS and Ethernet are used together, Ethernet is called to focus on only what it does best—point-to-point transport. MPLS, meanwhile, adds the connection-oriented capabilities that are a tool for creating SLA-backed services.

FIGURE 1

MPLS and Ethernet Come from Opposite Ends of the Network

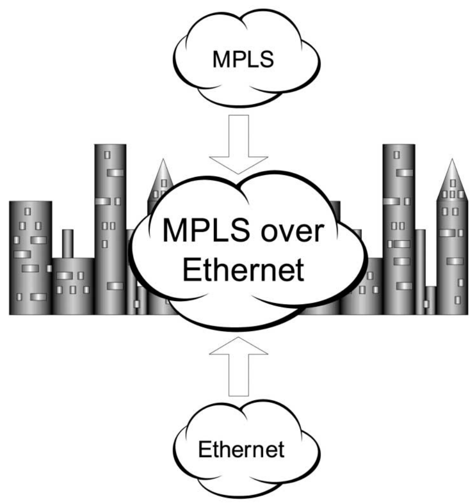

Creating SLA-Backed TLSs with MPLS over Ethernet

SLA Background

SLAs have come to attention with the increase of network outsourcing and the greater importance of inter-office networks. SLAs, from a customer's perspective, are crucial to optimal performance of the enterprise network and mission-critical applications. From the service-provider side, the ability to offer SLAs both sets the scope of possible customers and marks the critical point where network features actually translate into additional revenue.

The phrase "service-level agreement" refers to a contractual agreement between the service provider and customers, where the provider guarantees a certain level of service. Obviously, the guaranteed level of service can differ from customer to customer, as does the amount charged for such service levels, creating a differential pricing strategy for service providers. Traditionally, typical subjects of agreement are:

- Network availability (minimum guaranteed)
- Path/permanent virtual circuit availability
- Average round-trip network delay
- Effective throughput (network performance, burst allowance)
- Service-restoration time
- Billing methods (per usage or flat rate)

Over the last few years, SLAs have come to include more of a focus on application performance and other end-user metrics, such as the following:

- Application availability
- Application response time
- Performance in certain time windows

Many customers simply will not buy services that do not come with certain levels in a SLA. If a transparent LAN or externally hosted application is mission critical, one can expect a service offering without a certain level of SLA to be a non-starter. SLA capabilities are, therefore, a direct determinant of a service provider's potential customer base; hence the importance of premium SLA capabilities.

The Limitations of Using Ethernet to Create Transparent LAN Services

Ethernet is used in native form to offer TLS backed by an SLA, but the result is less than ideal. Most of the problems arise because Ethernet is a Layer-2 technology, ideally suited to handle point-to-point connections or LAN broadcast. Although carriers can currently offer SLAs over Ethernet networks, these SLAs have significant shortcomings, including lack of bandwidth reservation, slow recovery from network failure, and scalability limits. To illustrate, we'll use the example of TLS using 802.1Q-based virtual local area networks (VLAN) (see *Figure 2*).

First, it is difficult for the service provider to offer a guarantee of network throughput using native Ethernet. As pictured here, the service provider has configured two VLANs for customers A and B that extend across the metro network. However, the bandwidth on the trunk (and the metro network) is shared, and it is typically oversubscribed (2:1 oversubscription is pictured here).

Nothing in native Ethernet or 802.1Q prevents customer A from acting as a bandwidth hog and ruining any potential bandwidth SLA with customer B. A service provider could implement link-level controls on customer A, or use 802.1p to create a higher priority for customer B's traffic. But neither of these approaches creates the end-to-end bandwidth reservation necessary to guarantee high levels of network throughput.

Second, native Ethernet provides less than ideal resiliency features. The gold standard in this area is SONET, with its guarantee of >50-millisecond (ms) recovery time from fiber cuts. This feature accounts for the ability of service providers to achieve five-9s reliability and sell SONET-based services at a premium.

FIGURE 2

802.1Q -Based TLS—Lack of Bandwidth Reservation Makes it Hard to Offer Network-Performance SLAs

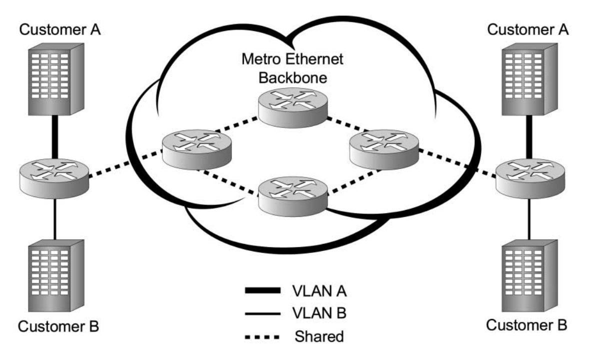

Ethernet relies on spanning trees to provide network resiliency. Unfortunately, spanning trees are not a good solution for MANs. They can take over 30 seconds to recover from an outage (either fiber cut or switch failure). Rapid spanning-tree technology is an improvement, but it still taps out at about a full second for reconfiguration.

Finally, the VLAN approach does not scale. From a service provider's perspective, 802.1Q's limit of 4,096 VLANs per network is very significant. It creates a serious upper limit on the metro TLS business model.

Finally, the Ethernet model comes within profound geographic limitations. The TLS service will be limited to the single network in question. A service provider, ideally, would like to offer SLAs not only within a single Ethernet network, but across WANs as well. This allows a service provider to generate revenue from customers connecting offices across the country.

Using MPLS over Ethernet to Offer SLA-Backed TLS

Using MPLS, a service provider has a powerful tool for offering premium SLAs in the metro. When MPLS capabilities are added to the TLS model, service providers can sell premium levels of service, such as bandwidth guarantees, fast recovery times, customized backup resources, or nation- or worldwide TLS (see *Figure 3*).

FIGURE 3

SLA-Backed TLS Using MPLS

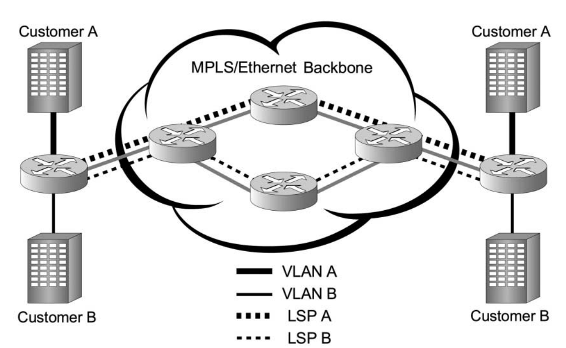

In the MPLS-based TLS implementation, each customer's VLAN is mapped onto an MPLS label switched path (LSP) that extends across the metro network. Each LSP enjoys reserved bandwidth across the MPLS cloud, as well as other QoS indicia (like jitter or latency limits). The LSPs created are circuits established by the IP signaling plane that can extend across an adjacent ATM network or an MPLS WAN without much fuss. This MPLS implementation allows the service provider to provide service levels critical to offering premium SLAs.

First, as just discussed, the service provider can now offer sophisticated, end-to-end network performance guarantees. The service provider can also offer a guaranteed amount of fixed bandwidth, as well as burst bandwidth contingent on availability. Note that this capability allows the service provider to oversubscribe with confidence, maximizing capital returns without jeopardizing the ability to offer bandwidth guarantees.

Second, service providers can use MPLS to offer both the speed of SONET recovery and the exact level of backup resources the customer is willing to pay for. The MPLS method for service resiliency generally depends on the establishment of an alternative, or recovery, path to complement the primary path. If the primary path is knocked out by a fiber cut or router failure, the alternative path takes over. The MPLS solution gives the service provider the ability to provision as few or as many resources for the recovery path as the customer desires. That way, providers can charge a premium rate to customers willing to pay for a backup path with as much bandwidth as the usual path. At the same time, providers can offer less expensive resiliency options for traffic that is important but not mission critical (see *Figure 4*).

Third, MPLS allows service providers to sell SLAs across a much broader range of networks. For example, in conjunction with an ATM or MPLS WAN, a service provider can offer cross-country or even international TLSs. Moreover, providers can accomplish cross-country and international services with the minimum of resource mapping. Ideally,

FIGURE 4

MPLS Service Recovery

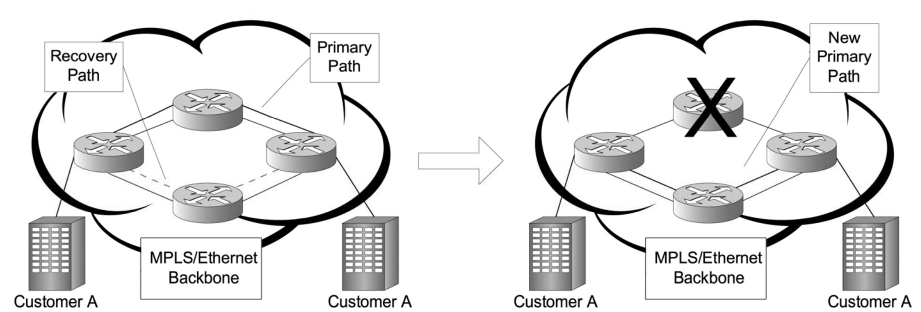

the service provider can provide a single, reserved path for coast-to-coast service (see *Figure 5*).

Finally, for providing TLSs, MPLS does not suffer the 4,096 limitation of 802.1Q. That means a service provider can sell SLAs without worrying about running out of VLAN space.

The Path to Long-Term Profitability

Let's see how the MPLS solutions discussed here can create a long-term profit model for service providers. The important fact is that the additional services need make only modest incremental increases in revenue to drive a large increase in profits. Consider a moderately successful mid-sized service provider earning an 8 percent operating margin on revenues of $250 million with a subscriber churn of 20 percent per year (see *Figure 6*).

Adding the MPLS services described can quickly lead to substantial gains. The provider can offer premium-level TLS and other services, along with nation- and worldwide connectivity, producing 3 percent growth in new subscribers. Targeting high-churn customer segments with customized, MPLS-based services may reduce churn by 3 percent. Finally, migrating existing customers to premium SLA packages can reasonably be expected to produce a 2 percent improvement in average revenue per subscriber. Collectively, these modest improvements will deliver an 8 percent boost in revenues. If we conservatively assume that the incremental service revenues carry a 50 percent operating margin, then the overall impact of this program will be a 50 percent improvement in operating profits.

Obviously, this is a simple model, and local conditions and operating costs will affect its viability. But it shows how a

FIGURE 5

SLAs Across the WAN

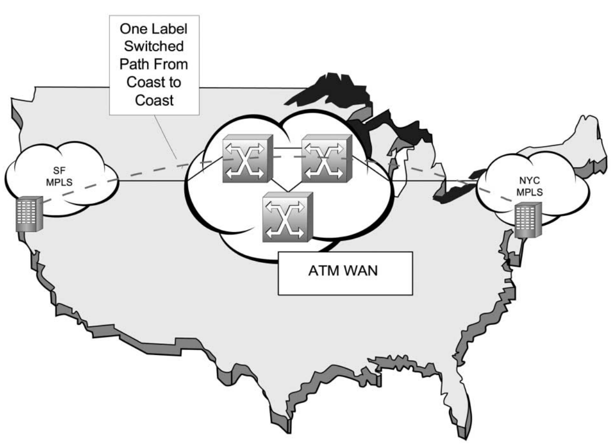

FIGURE 6

MPLS Services Driving Operating Profits

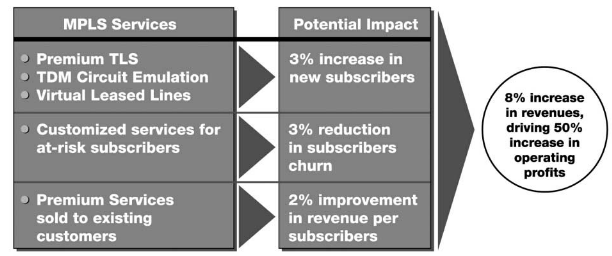

service provider can make more money on Ethernet networks than we might expect from a simple bandwidth-provisioning model.

Extensive experience with metro service providers shows that the following features loom large:

Interface Versatility
Ethernet isn't everywhere, yet, and for most metro areas, TDM connectivity remains critical to reaching a full range of customers. Packet over SONET (POS) remains important for uplink and SONET connectivity, and ATM can be useful in certain situations. The best routers support all.

Metro-optimized MPLS
As MPLS is deployed in metro areas, it is important to have MPLS implementations optimized for the features needed for MANs (essentially, an emphasis on service creation over traffic engineering).

Small-Form Factors
The simple economics of leasing and operational costs dictate routing equipment, at both the access ("basement router") and aggregation level, in a small-form factor.

Advanced Billing and Accounting Features
Eking the last penny of profit from metro networks turns out to be the key element of the profit story.

"Carrier-Class" Redundancy and Resiliency Features
Premium SLAs depend not only on the underlying capabilities of the network, but also on the resiliency of the equipment used. Look for routers with full switch-fabric and processor redundancy.

Conclusion

The "Ethernet revolution" is well underway. But it will take MPLS and similar technologies to add to the enthusiasm of a stronger profit model for service providers.

Cost-Effective Single-Fiber Broadband PON System

Muneer Zuhdi

Lightwave Technology Manager, Access Systems
Marconi Communications

Debra Wawro

Optical Component Manager, Access Systems
Marconi Communications

George BuAbbud

Vice President and Chief Scientist, Access Systems
Marconi Communications

Abstract

Recent advances in high-performance fiber-optic technologies combined with innovative optimization in the electronics have led to the evolution of a cost-effective single-fiber broadband passive optical network (PON).

This paper introduces a full-service PON system that delivers telephony, video, and high-speed data over a single fiber. A cost-effective method for the transport is presented. Simulations, analysis, and results are discussed.

Overview of PON Systems

PONs are widely used in access networks. The three major applications for PONs are fiber-to-the-cabinet (FTTCab), fiber-to-the-business (FTTB), and fiber-to-the-home (FTTH). An example of using PON in FTTH is shown in *Figure 1*. Since there are no active components between the central office (CO) and the end user, network-maintenance costs and requirements are dramatically reduced.

Because of the optical splitter, the cost for every channel at the CO will be shared by N homes. Current systems have splitting ratios between four and 32. Services delivered to the house are voice, video, and data. Video includes both cable television (CATV) and direct broadcast satellite (DBS) signals.

If a single-fiber solution is considered, voice and data operate bidirectionally on the 1310 nm wavelength, and video operates on the 1550 nm wavelength. The DBS and CATV signals could be transmitted in separate wavelengths in the erbium-doped fiber amplifier (EDFA) window. The asynchronous transfer mode (ATM) traffic carrying voice and data downstream is continuous transmission in the form of binary digital time division multiplexed (TDM) signals, while the upstream traffic is in time division multiple access (TDMA) bursts. Data rate varies between systems, but for a typical FTTH deployment, the rate is 5 megabits per second (Mbps) per home.

Overcoming Reflection and NEXT

PON systems are attractive to deploy in access networks because they are cost-effective. However, the optical signal has to accommodate high attenuation resulting from the splitting loss and the relatively long span. The splitting ratio and the distance between the optical line termination (OLT) and the optical network termination (ONT) are dictated by the sensitivity of the receiver. Reflection from the network or near-end crosstalk (NEXT) from the used devices could severely reduce the sensitivity of the receiver.

This problem is currently addressed by using a second fiber or a different wavelength to get good isolation between the transmitted and the received signals. A more cost-effective method is to shift the spectrum to higher frequency to get enough isolation to overcome the worst expected reflection, as seen in *Figure 2*, where we modulated the 25 Mbps signal on a 100 Mbps Manchester coded carrier.

We thus achieve an isolation between the two signals of 24.5 decibels (dB). However, instead of using 100 Mbps Manchester, we can use different line code with more concentrated spectrum to get better isolation. The line code also needs to induce transitions to be able to recover the clock and data quickly and with a minimum number of preamble data.

Simulations and Results

Figures 3 and *4* show the simulation results for the spectrum of the received signals at either end, along with the reflected signal from the transmitter of each end. The worst case of –14 dB of near-end reflection was considered in this simulation. The bottom chart shows that a minimum of 35 dB of isolation was obtained by shifting the spectrum of the

FIGURE 1

PON Architecture for FTTH

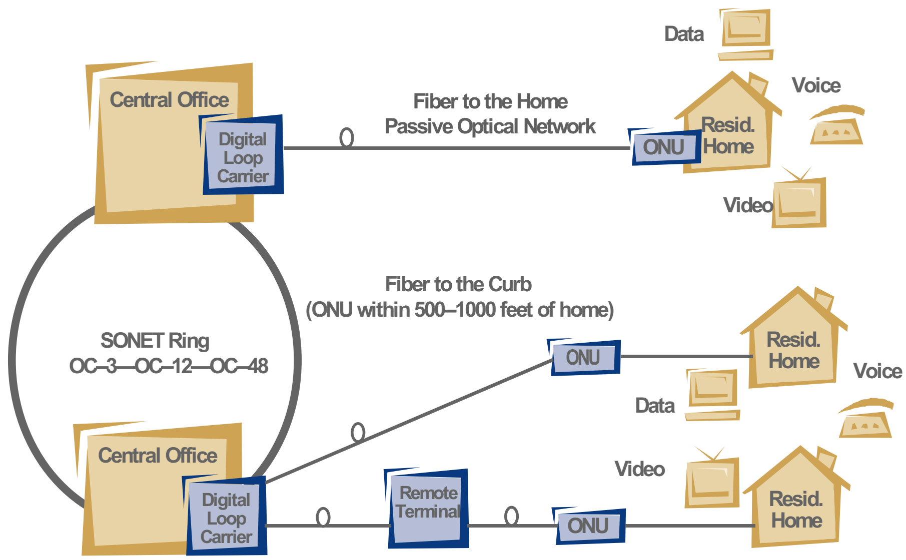

upstream signal. Although the signal on either direction could be shifted, it is better to shift the signal in the upstream direction. The additional cost in the receiver to recover the higher-frequency signal is shared by N homes, where N is the splitting ratio.

A PON system with splitting ratio of eight was built, simulated, and tested. A maximum attenuation of 17 dB was introduced in the signal path. At the same time, a reflection of –14 dB from the transmitter was inserted at each receiver. At both ends, simulations show that data could be recovered accurately at both receivers.

Figure 5 shows the eye pattern for the 25 Mbps and 100 Mbps signals when there is no reflection. *Figure 6* shows the eye pattern for the 25 Mbps and 100 Mbps signals in the presence of attenuation and reflection.

RF Video Transport

A FTTH solution based on wavelength division multiplexing (WDM) allows a variety of services to be combined on a single fiber. This flexible architecture enables the cost-effective transport of video services such as broadcast CATV, DBS, as well as two-way video-based services to the end user.

The video transport originates in the system from a satellite feed or CATV head end as a radio frequency (RF) signal. These broadcast video services are optically transported to the end user by subcarrier modulation (SCM) of two 1550 nm band lasers: DBS at 950 – 2050 megahertz (MHz) and CATV at 50 – 870 MHz. The relative power ratio and respective optical depths of the modulation of the two optical carriers are optimized to deliver a CNR > 48 dB and 25 dB for the CATV and DBS signals, respectively. The SCM modulated video signals

FIGURE 2

The power spectrum density for a NRZ signal is:

$$\mathrm{NRZ}(f) := \left(\frac{\sin(\pi \cdot f \cdot T)}{\pi \cdot f \cdot T}\right)^2$$

For a 25M signal, very good isolation can be acheived if the signal in the opposite direction is shifted to 100M. The PSD for a 100M Manchester signal is:

$$\mathrm{MCH100}(f) := \left[\frac{\sin\left(\pi \cdot f \cdot \frac{T}{8}\right)^4}{\left(\pi \cdot f \cdot \frac{T}{8}\right)^2 \cdot 4}\right]$$

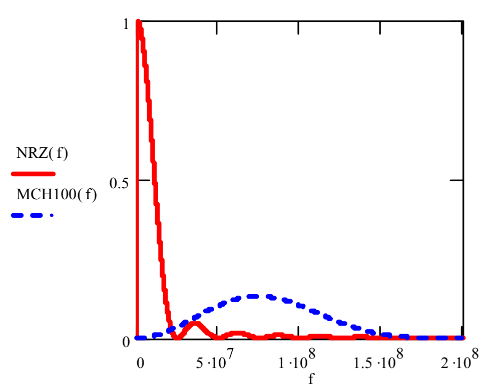

FIGURE 3

Isolation between 25 Mbps RX and 100 Mbps Reflection

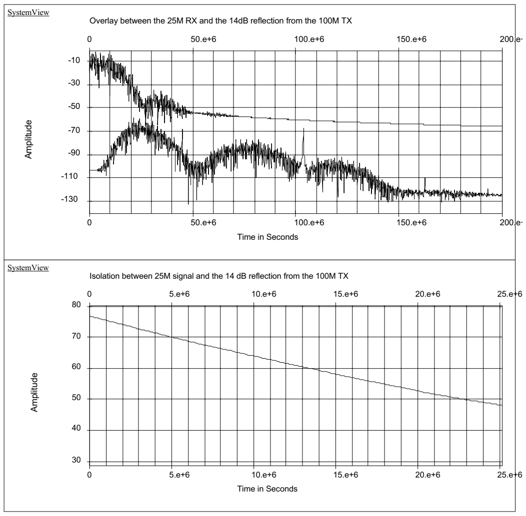

FIGURE 4

Isolation between 100 Mbps RX and 25 Mbps Reflection

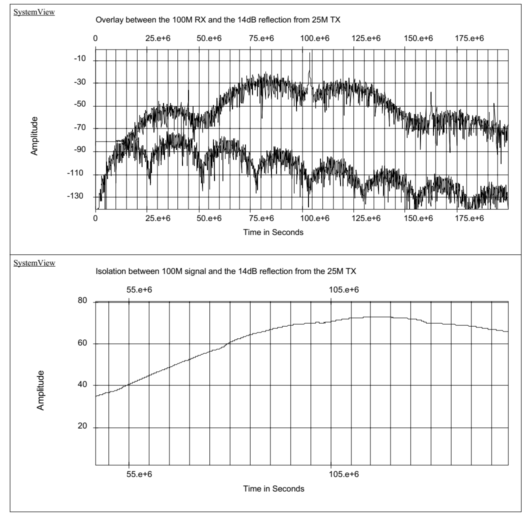

are distributed using low noise figure EDFAs. These EDFAs are optimized to provide high optical power over temperature extremes, while maintaining high carrier-to-noise ratios. The amplified video signals are broadcast and optically combined, as seen in *Figure 7(a)*, with a 1310 nm baseband signal for single-fiber transport using a splitter WDM cross-connect (SWX).

At the ONT, as seen in *Figures 7(b)* and *7(c)*, the video signals are separated from the baseband signal by optical filtering. Both 1550 nm window wavelengths are carried together to the video receiver photodiode and demultiplexed in the RF domain. The 1550mnm receiver sensitivity is maximized with a simple novel circuit configuration. In the home, there are standard interface-to-75 ohm coax ports from the RF video receiver.

Optical Components

The baseband digital signals are bidirectionally transported on a single fiber by using low-cost optical duplexers. In the OLT, a two-port duplexer houses a thermally stable 1310 nm laser diode (LD) and a PIN photodiode (PD) receiver. A beam splitter is used to direct the incoming 1310 nm signal to the receiver PD, while allowing the outbound 1310 nm signal to be launched on the fiber.

After the SWX, as seen in *Figure 7(a)*, the broadcast video signal is combined with the bidirectional baseband signal and is transported to the end user via a passive optical splitter (POS). Split ratios range from two to 32 homes and are typically placed within 30 kft of the CO without using any active devices in the network. The POS is designed to operate effectively in an outside plant environment, requiring a wide thermal range of –40˚C to +85˚C. Low-cost optical coupler technologies, such as fused biconic fiber couplers, and planar waveguides are adequate for this application.

At the home, the incoming signals are separated in the ONT by a three-port optical device, or triplexer. The triplexer separates the video wavelengths from the 1310 nm baseband signal by WDM optical filtering. The filtering technology allows better than 40 dB of isolation between the video and baseband

FIGURE 5

Eye Pattern for the Recovered 25 Mbps and 100 Mbps Signals

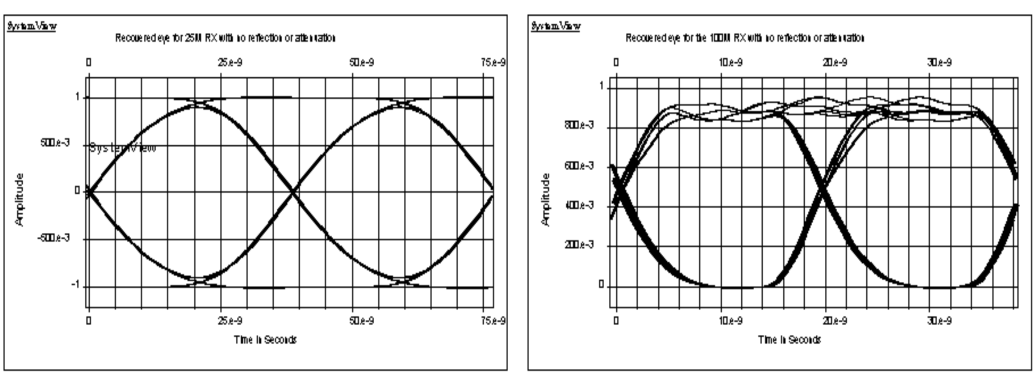

FIGURE 6

Eye Pattern for the Recovered 25 Mbps and 100 Mbps Signals in the Presence of –14 dB of Reflection and –17 dB of Attenuation

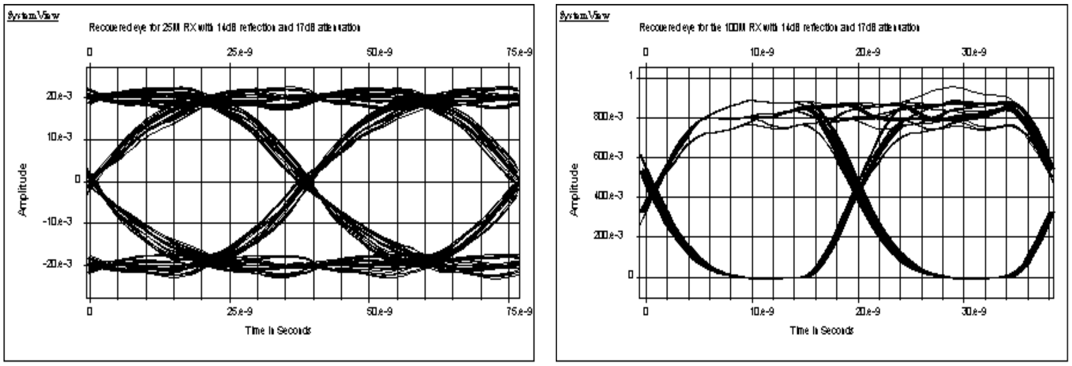

signals. The separated inbound optical signals are converted to electrical signals with a 1310 nm PD receiver and a 1550n m band PD receiver. The upstream baseband signal is launched on the fiber with a 1310 nm LD. The triplexer devices are thermally designed to operate in a residential garage or outside environment. The ONT converts the incoming signals into services required by the end user, including standard interfaces for telephony, high-speed data, and video services.

Conclusion

A cost-effective single-fiber PON system was presented. A new way of isolating upstream and downstream signals in PON systems was introduced and found to enable digital transmission on the same wavelength without limiting the sensitivity of the receivers. Spectrum shifting is more cost-effective than using a third wavelength or using a second fiber. A minimum of 35 dB of isolation was achieved. Good eye pattern for the recovered data at both receivers was obtained in the presence of worst-case reflection and worst-case attenuation. At the same time, we were able to transmit both DBS and CATV in the 1550 nm window without doing any WDM filtering at the receiver. These methods of transporting data and video enable us to provide a full-service cost-effective single-fiber PON system for FTTH application.

FIGURE 7

(a) Bidirectional transport of baseband signal is combined with broadcast video signals using SWX; (b) WDM filtering is used to optically separate the video and baseband signals in the ONT; (c) RF spectrum denoting allocated frequency bands

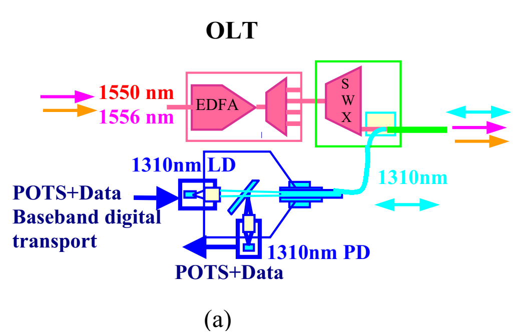

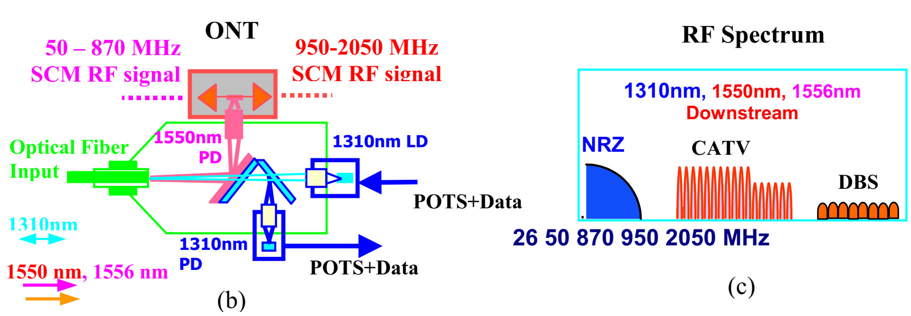

Section IV:

Packet Networking, IP Telephony, and the Internet

Leveraging VoIP to Enhance Competitive Positioning

Ron Aharoni

Director of Marketing, Advanced Testing Solutions Business Unit

NetTest

Trends

Voice over Internet protocol (VoIP) is a subset of next-generation network (NGN) topology, in which voice, video, and data "converge" into a single Internet protocol (IP)–based infrastructure, operating over both wired and wireless physical media. NGNs are associated with the transformation of current network architecture toward a multiservice and ubiquitous infrastructure, with unified control and management systems. Key applications include desktop voice and videoconferencing, audio and video streaming, and unified messaging—all of which are seen as potential sources of revenue for carriers and service providers. *Figure 1* depicts the exponential growth of data traffic, which is the driver behind the need for a data-oriented core networks.

There are two main reasons for this trend: 1) to eliminate the higher cost of operating and maintaining two separate networks and 2) to support new multimedia applications, which typically allow users to talk and exchange data and images in one single session. The technology has evolved into carrier-class gateways that function as the interface between IP–based networks and the public switched telephone network (PSTN), complete with signaling system 7 (SS7) support. Two key factors have contributed to the growth of VoIP. The first is the explosive growth of the Internet and other IP–based networks; the second is the trend to merge voice networks with data networks, resulting in lower operational costs.

Traditional voice links the signaling device directly to the Class-5 switch in the central office (CO), which presents it with a dial tone, collects the dialed digits, and interprets the user's request for service as a local call, long-distance call, special feature request, and so on. Under NGN architecture, the user's service request does not go directly to the Class-5 switch, but rather to the service agent, which is described as a "softswitch." The softswitch first determines the nature of the request, and then either routes it to a packet-switching network or to the conventional voice network managed by the long-established SS7. Thus, new and exciting voice features can be created without disrupting the equipment, practices, or stability of the current voice network.

The move toward the new architecture is being driven by the following:

- *Advanced solutions* enabling new ways to organize the way we work
- *Technical advances* that make it possible to implement converged solutions
- *Cost reduction* in installation, operation, maintenance, and use of the network services
- *Deregulation* allowing carriers to exploit the new developments and to offer competitive services
- *Industry standards* for open systems and integrated solutions

The stimulus for network change has traditionally been primarily cost savings. With the advent of the Internet and online access to millions of potential customers, however, supporting network-based applications that permit new ways of organizing and running the business has become the goal.

Applications

VoIP applications include the following:

- *Toll Bypass*: Routing long-distance calls through IP–based networks can reduce the cost significantly, especially for international calls.
- *Call-Center Integration*: Call centers usually use one voice link to talk to customers and one data link to retrieve information from databases. Using VoIP, these two links can be combined into a single cost-saving link. Furthermore, customers can go to the Web to find information and make calls to the call center from the Web page (click-to-talk).
- *Unified Messaging*: This application allows users to receive/retrieve various forms of messaging (voice, e-mail, fax, etc.) from a single access point.
- *IP Videoconferencing*: This application can help businesses reduce the costs of conferencing.
- *Corporate Intranets*: Many companies today have two separate networks: one for voice and one for data. These two networks can be integrated into one using IP telephony, reducing infrastructure and administration costs.

FIGURE 1

The Exponential Growth in Data Traffic Worldwide

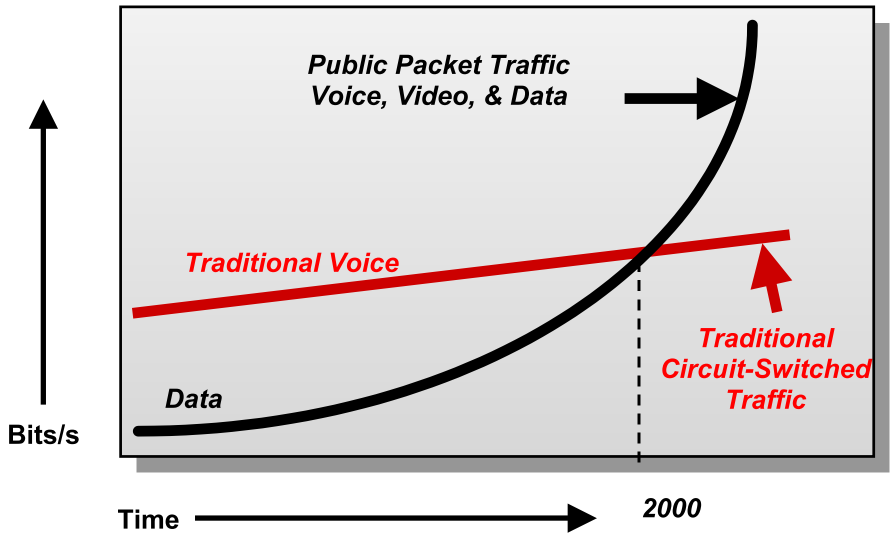

The Evolving Standards

Data transmission and voice transmission over packet-based networks have a major difference in requirements: Data is inherently loss-sensitive and delay-tolerant, while voice is loss-tolerant and delay-sensitive. For this reason, the transport layer in the VoIP protocol stack uses the user datagram protocol (UDP) to carry voice instead of the transmission control protocol (TCP). TCP is used to carry signaling messages such as call set-up and teardown. In addition, because most VoIP applications are real-time, the real-time transport protocol (RTP) is run on top of UDP to provide end-to-end delivery services for data with real-time characteristics. These services include payload-type identification, sequence numbering, time stamping, and delivery monitoring.

One of the common VoIP standards is the International Telecommunication Union–Telecommunication Standardization Sector (ITU–T) H.323, which was designed for multimedia communications systems and is considered bulky for simpler voice applications. As a result, the session initiation protocol (SIP) was developed by the Internet Engineering Task Force (IETF) as an alternative protocol offering less complexity and more flexibility. Both H.323 and SIP are peer-to-peer protocols. These protocols have been developed with the notion of "smart terminals–dumb network," meaning the software on personal computers (PCs) and workstations handle most of the protocol features in conjunction with very few network gateways.

However, a different philosophy puts the burden on "smart networks–dumb terminals," meaning the network is responsible for handling all new VoIP features using a new class of devices. These devices are controlled by protocols such as media gateway control protocol (MGCP) and media gateway control (MEGACO)/H.248. MGCP assumes a call-control architecture, where the call-control intelligence is outside the gateways and handled by external call-control elements, and hence is a master/slave protocol. The MGCP is a combination of two earlier protocols: the simple gateway control protocol (SGCP) and the Internet protocol device control (IPDC).

Figure 2 illustrates the generic VoIP network diagram. MEGACO/H.248 is an industry-standard protocol for interfacing between external call agents called media gateway controllers (MGCs) and media gateways (MGs). The standard is the result of a collaborative effort between the IETF and ITU standards organizations and is expected to gain wide industry acceptance as the official standard for decomposed gateway architectures sanctioned by both organizations. It was derived from the MGCP. The MEGACO/H.248 is used in Internet telephony and broadband access products, including trunking gateways, voice over asynchronous transfer mode (VoATM) gateways, IP telephony and residential gateways (such as set-top boxes and any type of digital subscriber line [xDSL] devices), and other cable–IP telephony equipment.

MEGACO/H.248 allows the separation of control-plane signaling and voice packet media transmission in the IP network, similar to the way it has been done in the PSTN and SS7 network, where the call establishment and termination are done via SS7, while the actual voice conversation is carried over the PSTN.

The next-generation decomposed architecture for IP convergence enables a network hierarchy in which a single MGC controls a cluster of MGs. This distributed architectural model presents a highly scaleable solution coupled with reliability and effective bandwidth utilization. The MG terminates voice flow to perform packet to circuit transcoding. The signaling gateway (SG) executes signaling technology adaptation. The MGC performs signaling interworking, call handling, and MGC functions.

FIGURE 2

Next-Generation Network Conceptual Architecture

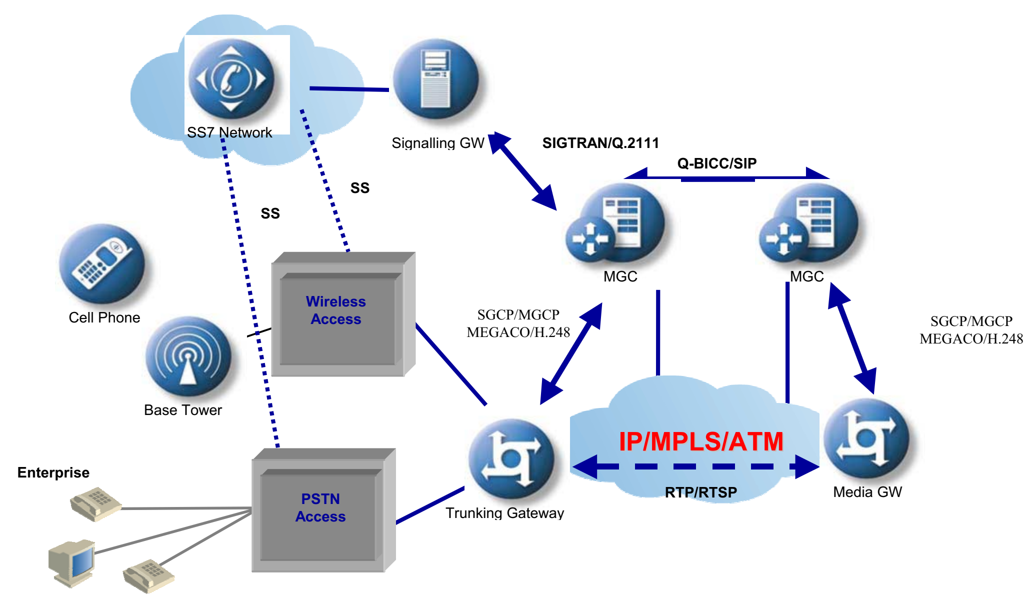

Technical Challenges

The existence of several standards creates the problem of interoperability. Gateways from different vendors may not be able to communicate with each other, not only because they support different standards, but also because many vendors have proprietary solutions that may yield better performance than standardized techniques. In addition, even gateways supporting H.323, for instance, may not work with each other since H.323 itself allows for flexibility in implementation.

Another important issue in VoIP is quality of service (QoS). The best-effort service available in IP networks today is not sufficient to provide VoIP service comparable to the traditional PSTN service in regards to reliability and voice quality. Hence, certain service-level guarantees are needed. So far, IETF standards such as resource reservation protocol (RSVP) and differentiated services (DiffServ) are promising in IP QoS support, and a number of VoIP gateway vendors have implemented them in their products. However, a significant amount of work still remains to be done in this area.

When implementing VoIP solutions, there are several technical challenges that we need to consider:

The Delay Factor
Excessive end-to-end delay makes conversation inconvenient and unnatural. Each component in the transmission path—sender, network, and receiver—adds delay. ITU-TG.114 (one-way transmission time) recommends 150 ms as the maximum desired one-way latency to achieve high-quality voice.

The Jitter Factor
Quantifies the effects of network delays on packet arrivals at the receiver. Packets transmitted at equal intervals from the left gateway arrive at the right gateway at irregular intervals. Excessive jitter makes speech choppy and difficult to understand. Jitter is calculated based on the inter-arrival time of successive packets. Jitter buffers (packet buffers that hold incoming packets for a specified amount of time) are used to reduce the effects of network fluctuations and create a smooth packet flow at the receiving end.

The Packet-Loss Factor
Typically occurs either in bursts or periodically due to a consistently congested network. Periodic loss in excess of 5 percent to 10 percent of all voice packets transmitted can degrade voice quality significantly. Occasional bursts of packet loss can also make conversation difficult.

The Sequence-Errors Factor
Congestion in packet-switched networks can cause packets to take different routes to reach the same destination. Packets may arrive out of order, resulting in garbled speech. Evaluating voice quality is a major priority for companies developing voice communication systems or products. To enable accurate, consistent rating of voice quality, the mean opinion score (MOS) was developed. The MOS uses the absolute category rating (ACR) procedure to determine the general acceptability or quality of voice communications systems or products. Under ACR procedures, evaluators are required to rate the overall quality of speech samples in support of telephone communications. Listeners use a five-category rating scale (i.e., excellent, good, fair, poor, and bad) with points assigned for each level, as follows:

5 – EXCELLENT
4 – GOOD
3 – FAIR
2 – POOR
1 – BAD

Following the evaluation, a numerical MOS score for overall speech quality is calculated.

Another method for analyzing the voice quality is the perceptual speech quality measure (PSQM). It is a means for objectively assessing the quality of speech that has been degraded by a telephony network. It has a high correlation to subjective quality across a range of types of distortion and is appropriate for testing networks that are subject to different coding types and transmission errors.

Defined by ITU–T recommendation P.861, PSQM is used primarily to test networks that have speech compression, digital speech interpolation, and packetization. Networks that carry VoIP, voice over frame relay (VoFR), and VoATM have these characteristics. However, the use of PSQM is not limited to these applications and can be used effectively to test wireless systems and cable-modem systems that carry speech.

Industry-Wide Focus on Areas of Growth and Contraction

VoIP presents service providers with new difficulties as relates to telephony, transmission testing, and POTS. Even in today's Internet environments, voice telephony continues to be a component that isn't going away. As these networks migrate to IP, the integration of voice services will become even very important.

To minimize the risk of deploying and/or integrating NGN devices in new or existing networks, carriers work with vendors on extensive pre-deployment testing. The advanced communications labs, which large carriers have, conduct conformance, performance, interoperability, and emulation of various scenarios. In conformance testing, a wide range of test procedures, which assess the adherence to industry standards, is conducted. Performance testing is aimed at stressing/loading a new device or a specific network configuration to assess its behavior under extreme loads. Interoperability testing is run on specific devices from the same manufacturer (or from several manufacturers) to ensure that once they are deployed and connected on the same network, they would work flawlessly. Similarly, emulation devices are connected to the new equipment, emulating specific protocols and/or standards, to assess the communications behavior of the new device prior and post deployment.

Using these procedures, carriers are gradually implementing NGN devices into their existing architecture. The extensive testing is used to ensure the QoS is maintained according to service-level agreements (SLAs) with their customers. The new technology has been implemented by carriers worldwide. According to RHK, Inc., many contracts to purchase and implement VoIP equipment have been signed in the last year. For example, Alcatel reported signing contracts for SGs with 26 customers, which include carriers from Germany, France, China, and South America. Ericsson has also reported NGN contracts with Germany, Chile, Spain, and South America. Nortel Networks, another international player, is reporting deals in Germany, Australia (Telstra), and other countries.

In *Figure 3*, we can see a forecast from RHK pointing to growing sales of NGN technology in the upcoming years. This forecast focuses on the North American market and covers a wide range of devices that comprise the NGN architecture.

Conclusions

Carriers are realizing that using packet-networks to carry all traffic types leads to lower operational costs. New revenues can also be generated due to the much higher flexibility in introducing new services to end users. These factors, coupled with the exponential growth in Internet traffic from new users and new applications, are strong indicators that the current trend toward NGNs will continue into the foreseeable future.

FIGURE 3

New Voice Infrastructure – North America

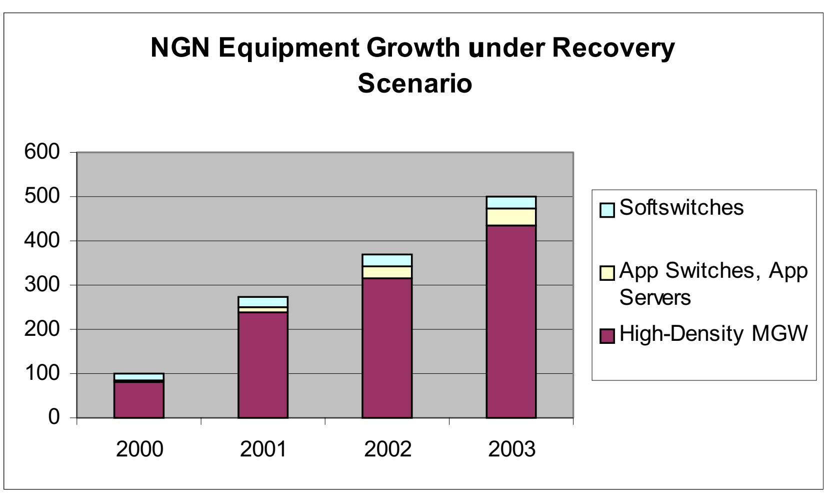

(Source: RHK, June 2001)

The Virtual Internet Data Center

Application Service Providers Evolve their Networks with Renewed Emphasis on Multiple Layers of Security

David J. Cavuto
Lead Engineer, VPN Firewall Team
Lucent Technologies

A Look Back

I remember the first time I called a Web hosting company. The representative didn't miss a beat—she had an answer for everything.

"Absolutely, sir. We take care of everything for you. We can make any changes you need. You just put your site up, and we make sure it stays up."

Then a thought hit me: How do I "put my site up"?

So I asked her. "Just how exactly do I put my site up?"

Seemed like a fair enough question. And she had an answer.

"Well, sir, you just create your files and transfer them to our servers using FTP."

I have to admit that I was a bit surprised. This company wanted me to transfer my files *across the Internet* to this Web server using file transfer protocol (FTP).

I tried to point out my concerns about using FTP to transfer files to my Web server. "But FTP sends passwords in the clear!"

This seemed to stump her.

"How do I connect to my database server, to create tables and upload data?"

"You can just use Microsoft Access to connect to the database."

What?! She was telling me that I could directly—over the public Internet—access my database server!

Things went downhill from there.

The story does have a happy ending, though. I was able to convince the site administrators to add some router access control lists (ACLs) to block would-be attackers, as well as create additional database and file system permissions to help limit the possible scope of intrusion.

I was surprised by the lack of knowledge—and lack of concern—that my provider had regarding security issues. After all, I could have some fairly sensitive information on my site.

E-commerce sites keep track of customer billing and payment data in databases. Web-based e-mail servers hold megabytes of private e-mail for individuals. Statistics servers hold usage logs and demographics reports of usage and user preferences.

The Problem Persists

In late 2001, a quick check of several Web-hosting companies revealed the ugly truth: Providers are still not overly concerned about security.

There is a strong tendency to hide behind SSL (secure sockets layer). *"Our site is secure—see the little lock icon!"* As a matter of fact, SSL only protects the data while it is *en route* (between your desktop and the server). What about how it is stored once it arrives? What protections do those systems offer against intrusion?

There have been several high-profile incidents lately regarding information stolen off servers at large companies. In one well-publicized event, the intruder stole credit-card numbers and "held them hostage," threatening publication unless he was paid a hefty fee.

While many individuals are fully capable of running hosting services, a good number are still not sensitive to security concerns, even when presented with overwhelming evidence that their users really *do* care.

Education is absolutely the best way to arm oneself against potential vulnerabilities.

But there actually is a good, quickly deployable solution to many security concerns. And it is available now, off-the-shelf.

Enter the VPN

VPNs, or virtual private networks, provide *cryptographically verifiable* security rather than just discrete network separation.

Even if a VPN–protected conversation is eavesdropped upon, the intruder can obtain no data contained in that conversation.

Three tenets of security are as follows:

- *Confidentiality:* Make sure no one can *see it*.
- *Integrity:* Make sure no one can *alter it*.
- *Authenticity:* Make sure we know *who you are*. And, as a distant fourth:
- *Non-Repudiation:* Prove that *only you* could have done it.

VPN technologies attempt to coalesce algorithms and procedures that provide varying degrees of each of these paradigms into standards-based protocols that will interoperate among different vendors. Unfortunately, there has been a profusion of such protocols. The list is fairly long:

- *IPSec* (Internet protocol [IP] security, part of the IPv6 standard but supported in IPv4 as well)
- *PPTP* (point-to-point tunneling protocol, used by Microsoft)
- *L2F* (Layer-2 forwarding, a Cisco Systems implementation)
- *L2TP* (Layer-2 tunneling protocol, an Internet Engineering Task Force [IETF]–standard mix of the previous two)
- And even our old friend *ATM* gets thrown into the mix
- As well as a latecomer—*MPLS*

The recommendations in this article, to a large degree, are written with IPSec VPNs in mind; other technologies may provide varying degrees of protection.

Each of these protocols functions differently, but the goal is the same: to allow messages to be exchanged over a public network, ensuring the aforementioned attributes *without requiring that the applications on either end of the conversation be security-aware.*

A Practical Example

Public servers, such as Web servers, must be accessible by the general public. Firewalls can be used to limit exposure to these hosts, but they can still be overtaken by traffic directed at transmission control protocol (TCP) ports allowed through the firewalls.

One way to think about this is that putting up a firewall is just like putting up a real brick wall. Nothing can get through, including valid user traffic. You have to open a port, which is like drilling a hole through the brick wall. An attacker can still put a gun up to the hole in the wall and do damage to whatever is exposed on the other side.

For most high-performance public servers, this is about the best you can do. To the extent that the server itself, the server operating system, and the server applications are kept up to date with patches and fixes with known security exploits, you have a good line of defense. Remember that security, like skiwear, is best when applied in layers.

However, some servers do not have to have general-purpose access to non-secure networks. For example, a database server, which is the core of many e-commerce sites, absolutely should not be directly accessible from the general Internet. But, of course, it must be reachable by the site administrators, as well as the Web server itself.

One way to achieve this segregation is to build a segregated network with a combination of VPN and firewall technologies.

Network Segregation and Risk Reduction

Public servers are protected from public attack by firewalls, preferably ones with protocol validation for the publicly accessible applications. Servers that don't need public access are protected by VPN gateways. A VPN gateway will only allow traffic to pass that has been properly decrypted and validated. However, since the public servers must access the private servers, non–VPN traffic is permitted to pass through the VPN gateway as long as it comes from the public servers.

End users are then given either a VPN client or another VPN gateway so they can access their own servers. A VPN client is a software agent that is installed on any computer that needs access to the protected resources. These non-public servers are now directly accessible by the end customer, but by no one else. This design allows end customers to be able to access their own private servers (like database servers) regardless of the inherent security provided by those servers.

It is important to remember that the public servers *can* be overtaken, even when you have followed the best-known practices of securing them. When (and not *if*) this occurs, the intruder will have access to whatever data on the *private* servers that the public servers would have access to—and without any restrictions imposed by the public servers' user interfaces (application service provider [ASP], common gateway interface [CGI], etc.)

So, it is absolutely critical to severely limit access initiated by the public server toward the private servers. Make use of any application-layer protection available—use read-only logins, database table views with limited permissions, file system and operating system (OS) privileges, SYN flood, and other flood throttling—anything to limit the amount of damage that could be caused.

The mindset to have is to assume that an attacker absolutely can overtake a public server and that they can gain access to *any* information on it, including raw script code, private server passwords, database login IDs, schema information, etc. Then ask the question: "What bad things could someone do with this information?"

And follow with the *critical* question: "What harm would this cause to my business?"

Virtual Data Center Components

The main firewall is the primary line of defense for the entire data center. This firewall must be engineered to defend against a wide variety of attacks and floods. It is completely self-defeating to install a security device at the *most critical point* in the network that cannot itself survive an all-out attack. If this firewall fails, all your customers are cut off. (And you had better head for the hills.)

Because firewalls are "stateful" devices, they are subject to differing types of floods and denial-of-service (DoS) attacks. The simplest attack is a "resource starvation" attack, where the intruder simply sends many packets in an attempt to cause the firewall to consume its entire available random access memory (RAM).

SYN flood, ACK flood, and fragment flood attacks are all different special cases of flood attacks, where the goal is to make firewalls or routers fail due to lack of either available memory, or central processing unit (CPU) cycles. Many firewalls react poorly when subject to some or all of these types of attacks.

It is critical that the main externally facing firewall be able to protect itself from these attacks as well as sufficiently insulate the servers that it is protecting. For example, if your firewall doesn't protect servers against fragmentation attacks, a Jolt2 attack may be able to compromise certain server operating systems.

Additionally, since there are some protocols that are intentionally allowed through, it may be desirous to perform some application-layer protocol validation.

In other words, if the firewall allows inbound connections to TCP destination port 80 (HTTP), it might be nice to make sure that the packets going to the server actually contain valid hypertext transfer protocol (HTTP) commands—and not random garbage designed to crash the server.

This type of inspection is typically very CPU–intensive and can be time-consuming as well. As a result, it is not commonly deployed on external links due to performance considerations.

However, it is extremely important to consider this type of inspection on the *internal* VPN gateway/firewall protecting the private servers. These connections are typically much less frequent than connections to the public servers (although may still be fairly high bandwidth).

But the information on those private servers is much more business-critical than information on the public servers. And applications services on the private servers may be less robust, from a security point of view. As a result, the risk associated with compromise of the private servers can certainly outweigh the potential performance degradation caused by application-layer protocol validation.

It may also be useful to perform some *content filtering* on traffic as it is entering the data center. For example, inbound mail can be virus scanned. Posts to news servers can be inspected for offensive language. Objects posted to Web servers can be scanned for hostile mobile code (Java/ActiveX).

Finally, *intrusion detection* should play a part on the internal data-center network. In particular, three classes of intrusion should be considered:

1. Evidence that a public server is under attack (e.g., many failed root login attempts)
2. Evidence that a public server has been compromised (e.g., a successful root login from the public network)
3. Evidence that a private server is under attack (e.g., port scanning of the private server address space)

Configuring and tuning an intrusion detection system to properly signal these conditions while minimizing the number of false negatives is sometimes difficult. Testing should be performed to ensure that the intrusion detection system would at least report positives.

Mutually Hostile Customers

Private data is protected from the big bad nasty Internet. Customers must connect via VPN to access their private resources. Individuals on the public Internet can only access public servers. All is right with the world, no?

No. Not yet. At this point, the service provider's customers are protected from the public at large. However, customers are NOT protected *from each other*. For instance, Customer A could overtake its own private hosts, then attempt to use them to attack other private hosts belonging to other customers.

The Need for VLANs in a Virtual Data Center

Virtual local-area network (VLAN) technology was originally designed to allow Layer-2 switches to be segregated, making one big switch into several small switches. Then it occurred to some folks that it would be nice to hook several such switches together, preserving the segregated boundary somehow, across a single link. Thus VLAN tagging—now known as IEEE 802.1Q—was born.

In a data-center environment, there is shared infrastructure (such as access links and edge routers) as well as individual customer infrastructure (usually hosts such as Web or database servers). The trick is to allow the individual infrastructure to connect to the shared infrastructure without compromising the security of individual components.

This can be accomplished by connecting the individual components to switches that can assign VLANs based on physical port (so the hosts don't have to support VLAN tagging). These switches then must aggregate at a security device that supports VLAN tags, so intra-customer traffic is filtered upon crossing a VLAN boundary.

This allows the security device (presumably a hybrid firewall/VPN gateway) to be shared across all customers, instead of requiring one per customer in the data center. This design assumes the existence of a firewall that supports 802.1Q VLAN tagged trunks, only a few of which exist.

In fact, if your security device can support virtual firewalls, all security functionality can be collapsed into a single choke point, providing firewall, VPN, and VLAN segregation.

VPN/VLAN Integration

The final step in this architecture is to ensure that the aforementioned VPN can be strongly associated with a VLAN. Once the customer connects via VPN, they must be confined to the VLAN associated with that customer and not allowed to go poking around at other customers' servers. The customers' VPN is extended into the VLAN. A *customer* is then equivalent to a *VLAN*.

Any private equipment assigned to that VLAN could be reached by that customer (or appropriate end users at that customer, anyway). More importantly, (from a security point of view), any private equipment *not* assigned to that VLAN couldn't be reached by any individuals at that customer.

From the service provider's point of view, this design is extremely scaleable. A single data center can grow to support thousands of different hosted customers, each with their own virtual data center. The cost of operations of this center is only marginally more than the cost of a supporting a single customer.

Most hardware—except dedicated customer hardware—can be managed for the entire data center, avoiding duplication, and allowing for economies of scale. Internet protocol (IP) addresses can be conserved because the "private" servers need only be reached through the VPN, which can use RFC 1918 reserved addresses.

And the customer sees a huge benefit and an immediate improvement in operations. Either they deploy software-based IPSec (VPN) clients to each PC that needs access to the "private" network, or they install a small VPN gateway on their network premises to allow them select access into their "private" space. Either way, their traffic is completely isolated from other customers' in the data center, giving them the impression that they have their own, completely dedicated network.

Some Final Notes

Enhancements can be made to additionally secure the data center. Shared servers can be multi-noded to provide different IP addresses for management versus public service. They can be multi-homed with two physical network interface cards (NICs) to provide an out-of-band management channel to ensure that no management traffic can enter via the public interface. Full redundancy (which was omitted in this exploration for the sake of simplicity) can be added. However, these are all variations on the theme.

The point I try to stress is that service providers need to become more sensitive to their customers' security concerns. While certain security concerns may seem a bit highbrow and paranoid, individuals are becoming more security savvy every day. The two most publicized technology-related events today are viruses and security intrusions, not fancy new servers or nifty virtual-reality software.

This is an ideal time for a data-center service provider to get into the security business, because it is still considered a value-added service. In a few years, you will either be secure, or you'll be out of business.

And in case you were wondering: Yes, there are service providers deploying networks in this fashion—right now. Customer service at these companies will also address these security concerns.

Architecting the Future of IP Services

Josh Cochin
Hardware Engineering Manager
Quarry Technologies

Steve Kohalmi
Chief Systems Architect
Quarry Technologies

1. Introduction

Network service providers (NSPs) are facing tough times. Lower than expected profits have forced many companies to re-examine their business plans. Offering basic connectivity and bandwidth services has become fiercely competitive, and margins on those services have suffered as a result. The ensuing price wars have taken their toll, and customer loyalty has been one of the main casualties.

Some forward-thinking providers have begun testing new, high-margin service offerings in hopes of differentiating themselves from the competition, attracting new customers, and retaining their existing clientele. The demand for advanced services, based on Internet protocol (IP) virtual private networks (VPNs) with application-level quality of service (QoS), data encryption, managed firewalls, and address translation, has not subsided, and many businesses have expressed a willingness to pay premiums for such special treatment of their important network traffic.

Unfortunately, to date, the trials of service-enabling systems have not gone as well as hoped. Though sound in concept, the first generation of IP service routers has exhibited some serious shortcomings when put to the test. These devices were designed for the era of 56 kilobits per second (kbps) dial access and over-subscribed residential digital subscriber line (DSL) connections. Most cannot handle the concurrent service mix and high-touch packet-processing throughput required by large numbers of higher-bandwidth business users who need to be sure their time-sensitive and mission-critical applications always receive appropriate precedence and security.

To solve this problem, a new architecture for the next-generation IP service edge switch is emerging. This innovative design uses high-speed application-specific integrated circuits (ASICs) in combination with programmable network processors to yield maximum packet-processing performance and feature flexibility. Interestingly, it is not focused on switching bandwidth or interface speeds, but instead on what really matters to support the business needs of service providers' customers—packet processing throughput per application flow, without any compromises.

2. No Compromises

As its name implies, the next-generation IP service edge switch is intended for use at the edge of the provider's network, where advanced services can be delivered most effectively. There, access links are aggregated and the traffic within those links is processed for optimal conduct over the backbone network. Logically then, the IP service edge switch sits behind any necessary access or aggregation devices and in front of the backbone routers. As such, it requires neither specialized access interfaces nor high-speed backbone ports, but simply the means to connect to them—technology aware, but not transport dependent.

By distinctly identifying subscriber traffic and classifying application flows at the network boundary, the IP service edge switch is able to apply QoS and encryption algorithms effectively and to aggregate flows prior to sending the traffic on to the shared IP backbone. Attempting to initiate QoS or security techniques elsewhere in the network could be futile, because the traffic would first have to cross the backbone in the clear where it might be subject to congestion and security breaches. The real challenge for the IP service edge switch is the ability to apply all required services concurrently to even the largest application flows, or class-based aggregate flows, without introducing performance-impacting latency or jitter. Additionally, because these services must be managed and ultimately billed for, detailed statistics collection is required. And, because the switch will be widely deployed in local communities, it must deliver all of these at a reasonable cost.

Such uncompromising performance is only possible if the IP service edge switch is built upon an architecture that is up to the challenge. Service providers looking to attract and retain lucrative business customers using advanced IP services will want to understand what makes these essential capabilities possible.

3. Architectural Issues

In general, an IP service edge switch receives data on its input interfaces, processes it, and, as part of this processing, determines the appropriate output interfaces, then sends it out those particular interfaces. Because the switch is designed for the edge of the network, this data-handling flow is not usually symmetrical.

On ingress to the network, many thousands of subscribers' data flows are aggregated onto a few backbone links. While the backbone links are quite large, it is still possible for them to be oversubscribed, and therefore it is important that subscriber traffic is properly aggregated according to relative priority of the type of the traffic.

In the opposite direction, data arriving from the network backbone has to be channeled into small subscriber access links. As diverse traffic flows are transmitted onto individual subscriber links that are much smaller in throughput than the core-side links, they have to be properly prioritized and controlled.

This aggregation of data flows requires a switch and packet processing architecture that can handle individual flows just as well as it handles aggregated flows and that can differentiate between different traffic types, properly manage traffic aggregation, and account for all data passing through the switch. Because the switch is deployed in a critical point of the network, it must also provide reliable operation even in the case of failures.

The basic components of an IP service edge switch are the input-output (I/O) module, the packet processing module, the traffic-management module, and the protocol-processing module. (Multiple, redundant instances of these components may be implemented to provide a high-degree of system reliability.) The protocol-processing module handles data traffic that is addressed to the system itself. Examples of this include routing protocol, system management, and tunnel-negotiation traffic. It is worth noting that data to and from the protocol-processing module requires traffic management as much as any user data going through the system to ensure smooth system operation and to prevent many different denial-of-service attack scenarios. However, because this module is not in the fast data-forwarding path, it will not be discussed here any further.

Many architectural choices must be made concerning how I/O, packet processing, and traffic manager modules are connected together. For reference, the most basic configuration would comprise one input, one packet processor, one traffic manager, and one output. A much more flexible and dependable system would have several of each of these components. But before such a system can be built, two issues must be settled.

First, the size of each component must be decided. Each I/O module obviously has to be sized in accordance with the particular media it services. Each packet-processing module has to be able, at a minimum, to handle all the processing requirements of the largest single traffic flow. Because the largest traffic flow can be as large as the largest I/O port in the switch, packet-processing bandwidth has to be sized accordingly. This means that if the largest I/O port is optical carrier (OC)–12, for example, the packet-processing module has to be able to process (terminate, decrypt, filter, translate, encrypt, and encapsulate) at this rate. Otherwise, the packet processor will back up and start dropping packets. This is unacceptable because the discard process is not intelligently managed based on QoS criteria, but rather is indiscriminately based on simple congestion.

Traffic of the same flow cannot be striped across multiple processors to increase packet-processing throughput because this could cause packets to be processed out of order, causing problems with tunnels that require sequence numbers (e.g., L2TP and IPsec) and many applications. Therefore, the packet-processing module must be sized for the task.

The second issue to be decided is how a traffic flow should be aggregated relative to all other flows. This aggregation decision must be made after packet processing has been completed. It is the foundation of the QoS treatment applied to a traffic flow. The traffic-manager module is responsible for this aggregation. Of course, it is possible to over-subscribe the switch such that more input traffic is received than the output ports can handle. But in a well-designed system, the resulting congestion would occur only at the output ports, where the traffic manager could handle it intelligently.

From the previous discussion, it should be clear that packet processing is an ingress function, while traffic management is an egress function. It should also be plain to see that in a perfect system all data destined to an output port would be received in the output port's traffic manager, where QoS and flow-aggregation decisions are made. The basic reference architecture therefore places the packet processor on the input side of the switch and the traffic manager on the output side (see *Figure 1*). When multiple I/O modules are connected together via an intermediate switch fabric, additional issues arise. If there are N I/O modules in the system, each with a maximum of P bandwidth, to prevent any chance of head-of-line blocking within the system, the switch fabric should provide 1xP input from each input module and NxP output to each output module's traffic manager. Thus, theoretically, the size of the switch fabric should equal NxNxP. In practice, this could be an unsustainable model due to exponential scalability requirements. However, it is intended to point out that the switch fabric must be sized to handle realistic traffic patterns between I/O modules and that an output of 1xP is clearly not enough.

Physically locating packet processing and traffic management together with the I/O module may seem like a good idea until the cost ramifications are examined (see *Figure 2*). There is a huge difference in the bandwidth requirements of I/O modules for an IP service edge switch ranging from low-speed access (e.g., T1, digital signal [DS]–3) to high-speed core-side interfaces (e.g., OC–12, Gigabit Ethernet). While one could use the same packet-processing and traffic-management hardware for all I/O modules, the cost to support lower-speed interfaces would prove to be prohibitively high. Conversely, optimizing the packet-processing and traffic-management designs for each I/O module type is not practical from the engineering resource and support aspects.

A better solution to the design of an IP service edge switch is to make packet processing and traffic management sepa-

FIGURE 1

General Data Flow in an IP Services Switch

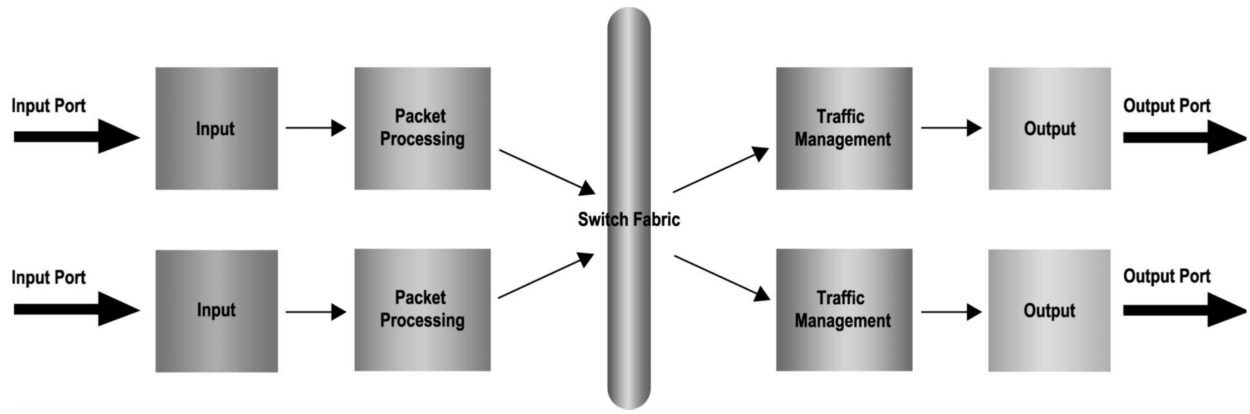

rate resources, which may be shared among multiple I/O modules (see *Figure 3*). This has the advantage of making efficient use of development and system resources, improving scalability, and facilitating system redundancy. Additionally, as seen, packet processing and traffic management are inherently separate processes (one being ingress and the other egress). As such, placing a switch fabric between them makes sense as well.

While the I/O, switch-fabric, and traffic-management modules are important components of the IP service edge switch, the real heart and brains of the system is in the packet-processing module. It is here that careful attention to design and optimization yield the greatest benefit.

FIGURE 2

System Architecture with All Resources on a Single Card

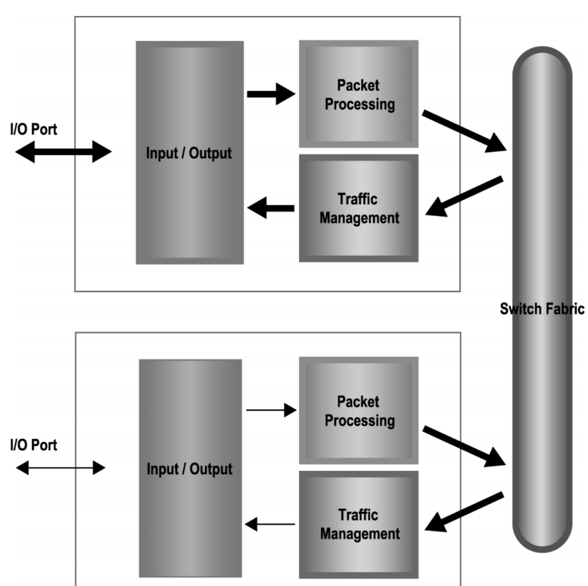

4. Packet-Processing Architecture

Processing packets is a multi-stage operation. It starts with the identification of the originator of the packets, and assigns this originator to a user group known as a subscriber. Then, databases related to this subscriber are opened according to information contained in the packet header. Tunnel termination is often required. This operation can be recursive because a tunnel may contain multiple subscribers, which may actually represent tunnels containing multiple subscribers, and so on.

Each subscriber may have its own service-level agreement (SLA) that requires the provider to deliver services and enforce policies relating to the data flow. Of course, the service provider, being in the business of making money, had better account for the services it is providing.

In addition to subscriber identification, there are many packet-processing stages that each packet must undergo. The packet-processing module performs filtering, traffic classification, metering, marking, address translation, routing, and encapsulation. In fact, several of these steps are performed twice per packet—once according to the ingress subscriber's SLA and again according to the egress subscriber's SLA (if the egress subscriber is local to the same switch). Statistics must be gathered for every packet throughout this processing to facilitate service management and billing.

The first-generation IP services routers employed general-purpose processors to handle traffic. The bandwidth of each of these individual processors is limited to around 30 to 100 megabits per second (Mbps) of packet throughput depending on the services performed, and data flow striping cannot be used to increase the throughput per flow due to packet sequencing issues (see *Figure 4*). Therefore, the two major problems with this approach should be immediately obvious: each processor is limited to much less than the potential bandwidth of an individual traffic flow; and even this meager throughput is dependent on the types and quantities of services that need to be performed. As intro-

FIGURE 3

System Architecture with Distributed I/O and Function Cards

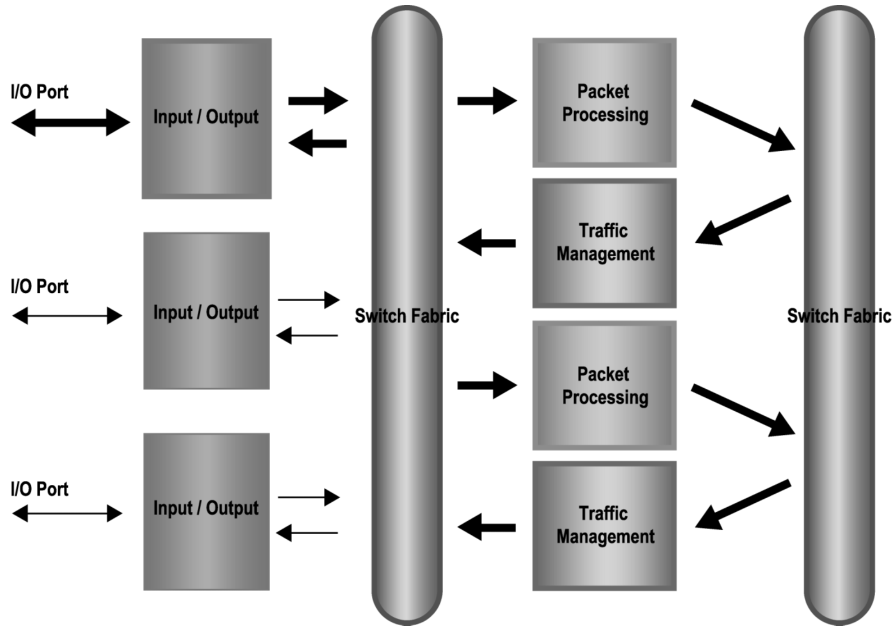

duced in the previous section, advanced IP services can be effectively applied only if neither of those restrictions is present. A packet-processing module must be able to handle the full bandwidth of the maximum size flow regardless of the processing complexity.

Pure ASIC–based solutions typically win the bandwidth race. Where they fail is in their product-execution risk factor and inflexibility. The service edge space is a fast moving, dynamic segment that continuously produces new standards and requirements. What a good IP service edge switch needs is the combination of the processors' programmability and the custom ASICs' speed.

Network processors came to the forefront in the past few years using multiple processing cores (e.g., MIPS, ARM) combined with specialized logic (e.g., CAMs) to provide a versatile approach to efficient data movement and packet processing. Network processors provide a step in the right direction but still cannot provide crucial elements of an IP service processor system, such as gigabit-speed encryption and deep, recursive traffic classification.

FIGURE 4

Data-Flow Striping and Merging Causes Packet-Sequencing Problems

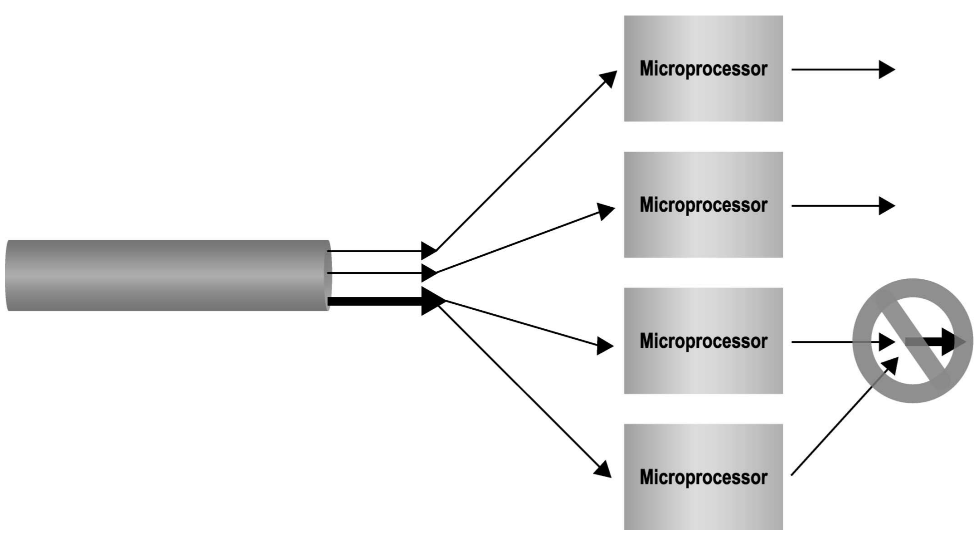

The best way to build a packet-processing module is to use a pipeline of network processors and assist them in the execution of special functions using ASICs and other specialty chips. In this way, the IP service edge switch is able to provide unmatched speed and flexibility while also avoiding sequencing issues of parallel-processing systems. Multiple pipelines can then be interconnected via the switch fabric to provide redundancy, load sharing, and system scalability.

The network processors can receive and transmit the full bandwidth of the largest port in the system, then distribute this traffic to ASICs and off-the-shelf coprocessors. Thus, the network-processor pipeline effectively provides sufficient bandwidth to and from these ASICs and coprocessors far exceeding the actual data throughput (see *Figure 5*). The ASICs and coprocessors provide specialty-processing functions. They are look-up, database, search, and encryption engines, providing adjunct processing to the network processors. All packet-manipulation decisions lie with the network processors. This division of labor improves system flexibility and allows for future growth and evolution of services.

5. Conclusion

Utilizing an IP service edge switch based on this next-generation architecture with its powerful packet-processing pipeline, NSPs will improve their ability to meet customer requirements and reap the associated revenues. Trials will prove the system's capacity to run multiple advanced services concurrently without compromises. Throughput per flow will be unrestricted so that QoS and security can be delivered at full line rate on all interfaces simultaneously.

Business telephone calls are confidential in nature and corporate financial transactions are usually secret. Likewise, voice traffic requires low latency, while transactional data needs high priority. The new IP service edge switch has the extraordinary capability to assure the privacy of this traffic even as it delivers the preferential treatment these crucial flows require. At the edge of the network, the switch aggregates the traffic streams by service class for safe, swift conduct over the backbone.

With network-based IP VPNs delivering highly desirable customized services, network providers can distinguish themselves as long-term, value-added business partners that contribute to the success of their customers. Offering such premium services will increase revenues and profit margins. Moreover, this will delight existing customers and attract new ones as the word spreads.

FIGURE 5

Network-Processor– and Coprocessor-Based Packet-Processing Pipeline

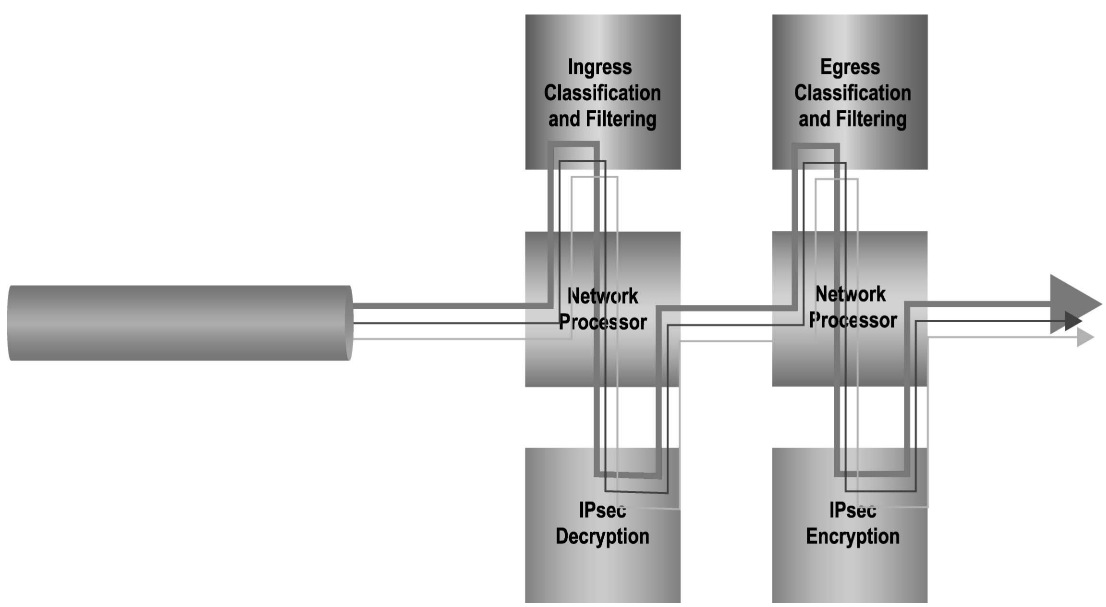

Sustained Profitability in a Packetized Voice Environment

Julia Cotton Sommerfeldt
Director, Product Management and Strategic Business Development
ADC Corporation

Introduction

During the past several years, there has been a gradual transformation of the public switched network from one that is based on time division multiplexing (TDM) to one that uses the Internet protocol (IP). As data traffic increased and network-management costs went up, the need for a converged network—one capable of delivering both voice and data services—became more apparent. Because IP is a pervasive protocol, the adoption rate for voice over IP (VoIP) accelerated. Proprietary Class-4 and Class-5 switches are now gradually being supplemented and replaced with a more standards-based, distributed architecture, which uses gateways and softswitches.

The reasons for migrating to VoIP and for offering other IP–based applications are compelling. However, to recognize the full benefits of a VoIP network, one must consider more than just the technology platform to deploy. This paper will outline the market drivers for IP–based voice networks, some of the expenses and complexities of deploying these networks, and the critical success factors from an operations standpoint that must be met to realize and sustain profitability.

Market Drivers

VoIP, voice over asynchronous transfer mode (VoATM), and voice over digital subscriber line (VoDSL) markets will grow at a rapid rate during the next five years (see *Figures 1* and 2). The largest growth of these three, according to many industry forecasts, will occur in the VoIP sector at a global level during the next few years. Service providers and corporations are looking at VoIP as the end game, because IP networks have the advantage of providing additional multimedia applications, such as unified messaging (UM) and video services. Network endpoints, such as telephones, personal computers (PCs), and personal digital assistants (PDAs) can be IP enabled, so service providers can deploy a variety of IP services on a mass scale. While ATM currently has the advantage in quality of service (QoS), the quality issues in VoIP are being addressed and will become less of an obstacle during the next year. Private IP networks, versus the public Internet, enable the means to manage QoS on an end-to-end basis.

The drivers of VoIP markets include the following:

- International deregulation and privatization of telecommunications industry
  - Increased competition in deregulated markets
  - VoIP is not currently regulated over public, private, or virtual private data networks.
  - Toll arbitrage (avoiding tariffs) enables service providers to address price sensitive markets.
- Growth of global IP networks
- Relative to circuit-switched networks, the time to market to build a VoIP network is quicker, and the costs are lower.
- Demand for advanced applications that use voice and other media, and intelligent IP–enabled endpoints.
- Increase in IP–based intelligent endpoints (e.g., PCs, PDAs, phones)
- IP is becoming the pervasive protocol from the core to the edge and out to the endpoints.

The growth of IP networks and intelligent IP endpoints set the stage for more advanced and integrated applications such as UM, voice-enabled Web pages, voice-recognition systems, videoconferencing, collaborative computing, and IP–based call centers. Mobility is a common, or additional, requirement across the applications spectrum.

Some of the applications and their time frames driving VoIP deployment are as follows:

2000: Toll arbitrage, prepaid calling cards, wholesale services (carrier's carrier), settlement and clearinghouse (for global reach), and Internet call waiting
2001: Toll arbitrage in lesser-developed countries, UM, interactive voice response (IVR), wireless/mobility applications, and voice-enabled Web sites
2002: E-commerce, call centers, collaborative computing, videoconferencing, voice-enabled Web sites, wireless/mobility applications, and increased VoIP deployment in the enterprise market
2003: Further advancement and integration of these systems

Customers have high expectations and demands of their service providers. There is a pent up demand from both consumer

FIGURE 1

Voice-over-Packet-Network Services—North America

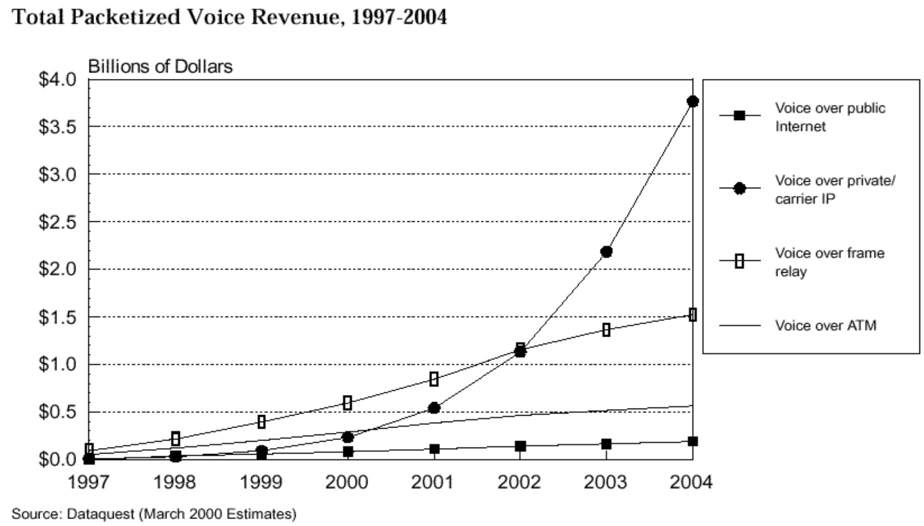

FIGURE 2

Voice-over-Packet-Network Services—Western Europe

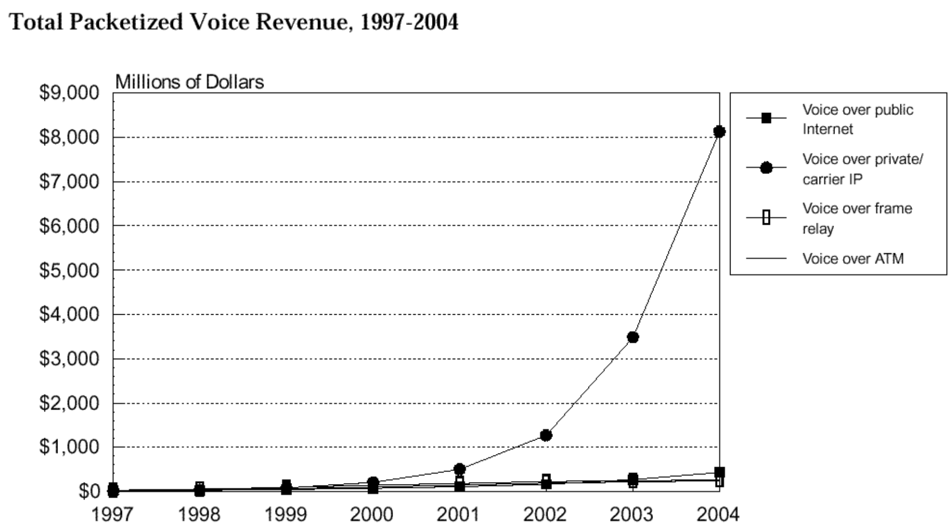

and business markets for advanced applications and services. Customers will continue to have high expectations that their providers will offer them these services quickly and reliably. The service providers, incumbents, and new entrants need to respond to this market push.

For service providers to keep up with this demand, their networks must be stable, reliable, secure, able to scale quickly, and easy to manage. Incumbent and next-generation service providers see the value of building and managing one network to deliver multiple services—such as voice, data, and video—versus managing separate infrastructures for those applications. More and more voice service providers are migrating from circuit-based networks to packet-based networks so they can deliver advanced applications at lower costs.

Reducing Expenses and Complexities

How do we lower these costs? The idea is that if we build one network capable of transporting voice and data applications, rather than two separate and distinct networks for voice and data, we will minimize expenses. The critical factor, which is often an afterthought, is that one can maximize those cost reductions only if the converged network has the appropriate, next-generation management tools in place. As with traditional TDM–based voice networks, cost reductions in VoIP networks are specifically related to network operations, administration, management, and provisioning (OAM&P), which can amount to 80 percent of the overall expense of owning and operating a network.

There are other important short-term and long-term factors that we need to consider regarding cost reductions. How do we integrate disparate new and legacy systems into a seamless solution? In the transition period, we are adding gateway equipment and software to bridge these two networks. As the number of network elements (NEs) increases, so does network complexity. In the short term, there are more elements to manage. During this migration period, the operations support systems (OSS) used must be capable of addressing legacy and next-generation technologies, elements, and services in an integrated manner.

The long-term consideration is the need for flexibility and growth management. The capability for growth must be inclusive of the network infrastructure, the applications, and the network support systems. In the long term, the intelligent, converged network will only be as intelligent as the end-to-end OSS supporting it.

In many cases, often in a reactive mode, some service providers tend to install separate support systems for each area of functionality. Soon, there is a plethora of disparate systems that are not interconnected. This only stresses the bottom line rather than improving it. Selling more services, especially new and bundled services, will increase revenues, but to turn these services up quickly, or "on the fly," and to increase the margins, an integrated OSS must be in place. Without that, the ratio of sales to expenses is not likely to change for the better. Writing customized interfaces for these is apt to take a long time and can lower one's margins. In this highly competitive marketplace, quick time to market and lowered operational costs are vital.

Although we can deliver voice and data applications using a common protocol (IP), the requirements encompassing those two applications differ. We also need to consider that the circuit-switched network technologies and the packetized voice technologies are generations apart. There is a clear need to integrate these technologies and the system that enables effective operations. We may expect that network technicians with data backgrounds will understand voice networks and vice versa. The technicians that are highly skilled, experienced, and knowledgeable in both network generations are limited in number and increasingly demand higher salaries. At the same time, we are reducing the ratio of employees to access lines (see *Figure 3*).

This indicates that we need to handle these complexities through automation. An intelligent OSS can handle some of these complex processes in the background so that less experienced employees can operate the system. This increases the labor pool, reduces training costs, and reduces recurring operating expenses.

Another way to reduce costs is to reduce customer churn. Because customer satisfaction and retention are key factors in this highly competitive landscape, service providers need the ability to quickly provision a variety of services while offering service-level agreements (SLAs). Having an OSS that is part homegrown and part off the shelf, and that is not integrated, does not enable one to quickly create and deploy services. This, in turn, negatively affects revenue potential.

Some functions in the OSS process include taking a service order, designing, activating, monitoring, trouble shooting, and billing for the service with an continual, complete, and accurate understanding of the NEs. An integrated, end-to-end solution enables the tracking of all orders and requests received so we can ensure that they are moving toward completion and not getting dropped somewhere in the process. This directly correlates to efficient, timely, and accurate customer service.

How do we know what is in our network? An inventory-management system can track what is in the network and its availability. When this system ties into a product and services catalog, we can more accurately create bundles and deploy them when a customer orders them. A system tracking these orders will appropriately and automatically decrement the inventory once it is deployed.

While one can have separate systems to manage each of these processes, the most value comes through leveraging the seamless integration of the systems. Integration is the

FIGURE 3

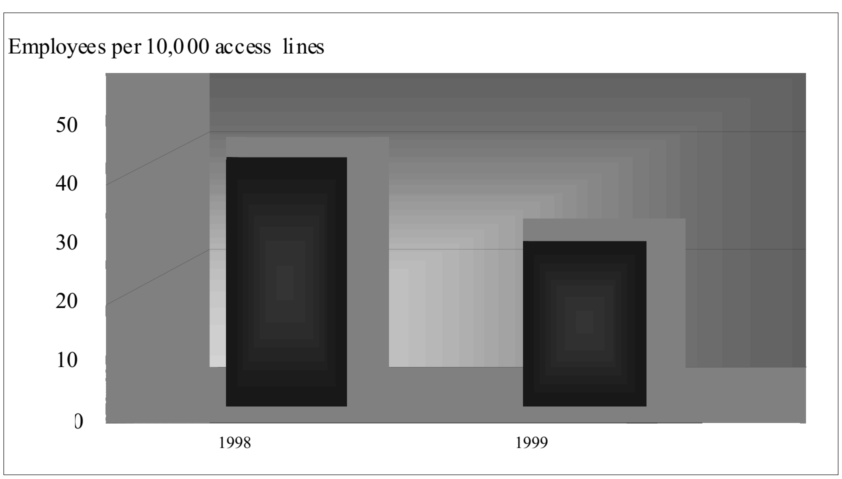

Source: CIBC Oppenheimer

key and can meet these challenges. A model where workflow provides an end-to-end business process automation to perform the myriad of tasks in an orderly manner is a necessity (see *Figure 4*).

In the next-generation environment, the service provider's objectives to minimize risks, integrate applications, provision on the fly, accurately charge and bill for services, increase flexibility, and empower customers and customer-service representatives (CSRs) with intelligence are heightened.

There will be a transition period to this new IP–based voice network, so the OSS must work in a multivendor environment and the OSS firm must fully understand your business. Modularity and flexibility in the network equipment and in the OSS is key to accommodating network growth (see *Figure 5*).

Conclusion

A migration to VoIP networks can significantly lower costs and help service providers to sustain competitive pressures. At the same time, an IP–based network opens the doors to new revenue streams from other multimedia IP–based applications. While the hype around this is tremendous and the market forecasts are alluring, one can only maximize the gain by deploying a complete and integrated OSS that is capable of managing both generations of a network. To minimize expenses and maximize profits, the OSS must support circuit-switched voice technologies, data technologies, and the emerging VoIP technologies. With network OAM&P costs accounting for as much as 80 percent of the overall expense of owning and managing a network, the implementation of an end-to-end OSS system cannot be an afterthought.

FIGURE 4

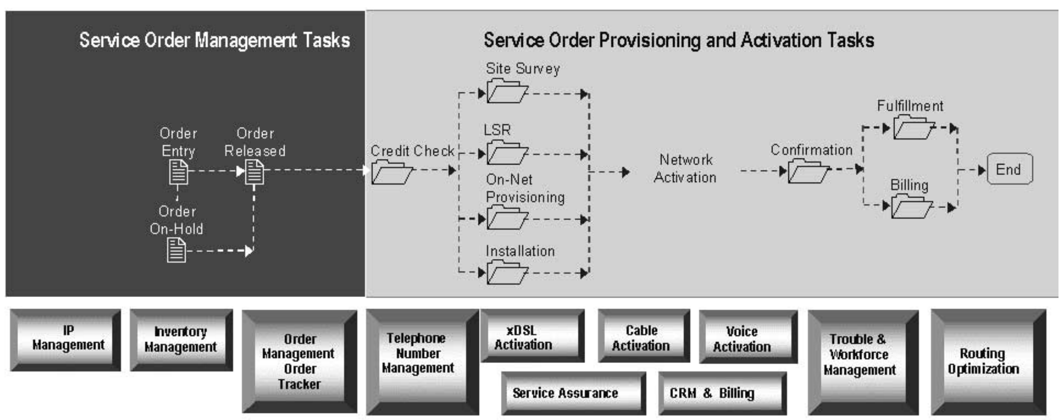

FIGURE 5

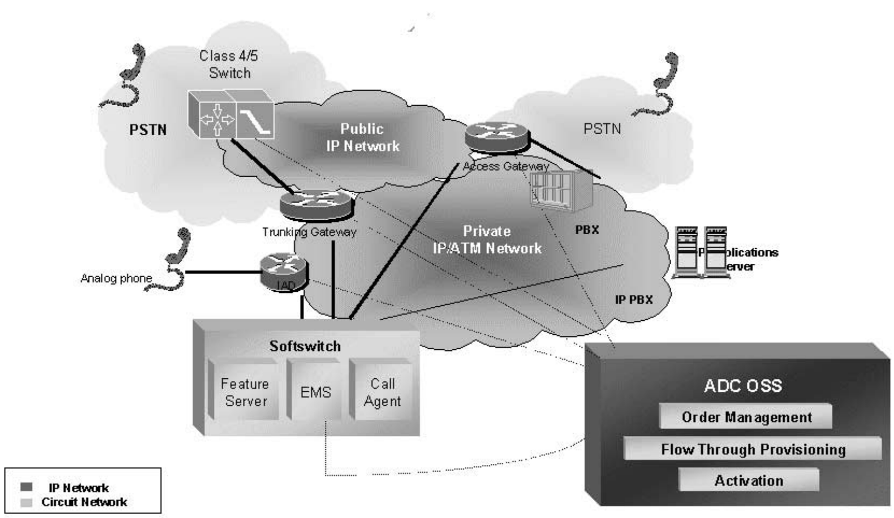

VoiceXML: Speech and Telephony Meet the Internet

Deborah A. Dahl
Manager, Advanced Development, Natural Language Speech Solutions Group
Unisys Corporation

Abstract

VoiceXML is an emerging standard for developing telephony speech-recognition applications. Because VoiceXML uses Internet technology, VoiceXML applications can easily leverage work that has been done for graphical Web applications. This paper describes VoiceXML, reviews its history, discusses its advantages and disadvantages, and talks about where VoiceXML may go in the future.

Introduction

Using the Internet both as a source of information and to perform e-commerce transactions has become astonishingly routine during the few years that the World Wide Web (WWW) has been in existence. People go to the Internet many times a day for countless reasons, and usage is increasing exponentially. But accessing the Internet through the most typical routes can become pretty cumbersome, especially when we are away from our desktop computers and wire-based networks. What if we could access the Internet through smaller, more ubiquitous, and, frequently, wireless devices—such as the telephone? Routine information—weather, stock quotes, news, sports scores, train and airline schedules—would be available anytime and anywhere. Transactions such as making reservations, banking, stock trading, and shopping could be done at one's convenience.

Clearly, there are obstacles to this vision, and an important one is the limited graphical bandwidth of handheld devices, such as telephones and personal digital assistants (PDAs), compared to desktop and laptop computers. To be handheld, these devices must be small, but small displays and keypads limit the extent to which information can be transferred visually. Although current mobile telephones have increasingly sophisticated display capabilities, any really practical telephony-based Internet access is going to have to involve speech output at a minimum and, given the limitations of keyboards and keypads on small devices, also frequently involve speech recognition.

Speech has long been recognized as the most natural way for people to communicate with machines, but machines that can recognize, understand, and act on human speech have been much harder to create than anyone ever imagined when early speech-recognition research began in the 1950s. Although we are still many years away from computers with speech understanding capabilities approaching those of humans, simple speech-recognition applications are starting to become commonplace (a report from Jupiter Media Metrix dated September 7, 2001, states that 18 percent of consumers between 18 and 34 had used speech-recognition technology to make a purchase). Even the simplest speech-recognition applications have the potential for significant cost savings (it is claimed that automated voice systems in call centers can reduce the cost of each customer transaction to $1 to $2, compared to $40 to 60 for calls handled by a live agent [1]). Speech applications can also greatly improve user satisfaction as compared to telephone applications that rely on the limited interface provided by touchtones.

Fortunately, speech-recognition technology has come far in the last few years, and current speech-recognition technology is very usable for many applications. It is still true that interacting with current speech applications is very different from the experience of talking to a human. However, when they are well designed, these applications serve their purposes effectively—speech recognition is acceptable to most users, prompts are easy to understand, and dialogs are designed so that users can focus on their tasks, not the speech interface.

The first applications of speech recognition in telephony were based on integrating the technology with commercial interactive voice response (IVR) systems. These systems traditionally use touchtones to gather user input, using proprietary approaches. This strategy led to effective but expensive systems that required the direct participation of speech-recognition experts in the development process. A new approach was needed to make developing speech applications fast and cost-effective.

In March of 2000, a specification was released that describes a new approach to developing telephony-based speech applications using WWW technologies. The goal of this specification, VoiceXML 1.0, is to enable users to access information, particularly Internet information, using speech on small devices, especially the telephone.

History

VoiceXML is the result of the convergence of three technologies: the WWW, computer telephony, and speech recognition. In the mid-1990s, researchers at several companies began to realize the potential for synergy among these technologies, and they began developing ways for allowing telephony-based speech applications to integrate with the Internet. These early companies soon came to understand that a standardized way of supporting telephony-based speech applications would be far more potent than continuing to develop multiple competing approaches. To move toward a single standard for this technology, Lucent Technologies, AT&T, IBM, and Motorola joined together to form the VoiceXML Forum [2], which released, in March of 2000, the VoiceXML 1.0 standard. While the VoiceXML Forum was designing VoiceXML 1.0, the World Wide Web Consortium (W3C) [3] was also addressing the question of integrating speech and the WWW. The W3C chartered the Voice Browser Working Group [4] in February of 1999 to define requirements and standards relating to spoken dialogs and the Internet. The VoiceXML Forum submitted VoiceXML1.0 to the W3C in March of 2000 for consideration as the basis of a standard for defining spoken dialogs. The Voice Browser Working Group has been working on a new version of VoiceXML, VoiceXML 2.0, since then.

What Is VoiceXML?

VoiceXML (voice extensible markup language) is an XML–based [5] markup language that defines a spoken dialog just as hypertext markup language (HTML) defines a graphical Web page. VoiceXML has generated tremendous interest among the speech recognition, telephony, and Internet industries. The fact that there are currently 566 companies that are members of the VoiceXML Forum is just one measure of this interest. Let's see what some of the reasons for this excitement are.

VoiceXML lets a developer define what the computer is going to say to the user, what the computer expects the user to say in return, and what to do with the information provided by the user as well as telephony operations such as transferring a call or hanging up. Proprietary IVR platforms have these capabilities as well, so why would an enterprise that is considering deploying a speech application choose VoiceXML over proprietary IVR systems?

VoiceXML has many advantages over proprietary systems:

1. ***Open Platform:*** Proprietary platforms are just that, proprietary. Early speech applications were developed using proprietary IVR tools and run on proprietary IVR platforms. Changing platforms is extremely labor-intensive—dialogs have to be rescripted, speech grammars need to be rewritten, back-end integration needs to be redone. Application developers need specialized training for every IVR with which they work. Every platform makes use of specialized tools that are expensive to develop and time-consuming for developers to learn. VoiceXML provides an open environment with standardized dialog scripting and speech grammar formats that developers have to learn only once. Furthermore, because it is based on XML, a vast selection of generic XML editing and parsing tools is available, including both commercial and freely available open-source tools.
2. ***Leveraging Existing Web Integration:*** Most applications also require at least some level of back-end integration with databases hosted on enterprise servers. Performing this integration is a labor-intensive process with little potential for reuse of the integration software. VoiceXML allows this back-end development to be done once for any Web environment, whether spoken or graphical. In addition, back-end development is done with standard Web development techniques for both client and server-side development, such as CGI scripts, Java servlets, JSP, and ASP. This allows application developers to leverage existing Web development skills rather than learning proprietary tools.
3. ***Declarative Representation of Dialogue:*** VoiceXML attempts to make the description of the dialog *declarative*. This means that the developer describes the dialog that is to be conducted at a relatively high level. The Voice Browser platform handles the low-level details of transitioning between dialog states, activating speech recognizers, handling telephony, and playing prompts. In contrast, most current proprietary application development environments are *procedural*, in that they require the developer to handle these details manually, through programming in languages such as C++, Java, Visual Basic, and proprietary IVR scripting languages. In practice, real VoiceXML applications do require some procedural programming on the part of the developer, but VoiceXML attempts to minimize it. Of course, high-level programming environments always come at the cost of some loss of low-level control, and it is not clear that VoiceXML has achieved the right balance between these goals. However, in time, this balance will get sorted out as the industry gains experience with many speech applications and we gain a more complete understanding of how much low-level control is really necessary to good applications.
4. ***Ease of Learning:*** VoiceXML is easy to learn compared to other approaches, such as proprietary languages or full programming languages, such as C/C++ or Java.
5. ***Open Standard:*** VoiceXML is an open standard that has many advantages:
   - Because VoiceXML is an open standard, a company's investment in its VoiceXML applications is not vulnerable to the business decisions and stability of a single vendor, which might go out of business, fail to deliver promised features, or discontinue support of a proprietary product.
   - Through the W3C process, any interested company can have input into the direction of the standard by joining the W3C and participating in the Voice Browser Working Group.
   - VoiceXML is being exercised by more developers (by orders of magnitude) than any proprietary language, making it much easier to understand what features and potential features are truly important and which are peripheral. (One platform alone, VoiceGenie, claims to have 2,950 developers in 90 countries in their developers' network.)

Although in many ways VoiceXML is at this point less powerful than proprietary approaches, it exhibits all the hall-marks of a disruptive technology in the sense used by Clayton Christensen in *The Innovator's Dilemma* [6]. It is cheaper and initially less powerful than existing technology. In addition, it may be viewed with disinterest by key large customers whose chief objective is getting new features added to their current implementations and who are not interested in newer but less powerful approaches. As Christensen points out, however, the failure of a company to correctly respond to disruptive technologies invariably has severe negative consequences when the technology matures and customers start demanding it. The tremendous interest in VoiceXML will create a great deal of pressure to drive the technology to maturity very quickly. Consequently, companies that have a vested interest in proprietary systems for speech and telephony integration need to be involved in VoiceXML, or they will be left behind.

Example: Defining a Spoken Dialog in VoiceXML

One attractive feature of VoiceXML is that there is not a lot of overhead for simple applications. Simple applications are simple to define. Of course, any practical application will be far more complex than the simple example w wi'll be looking at, but this initial simplicity is very helpful in reducing the learning curve for VoiceXML. It is actually possible to get something to happen in VoiceXML with very little effort. The availability of free voice browser implementations that are appropriate for learning VoiceXML and doing simple prototyping also helps to reduce the learning curve.

We can see many of the key features of VoiceXML), which asks the user what beverage he or she would like, in *Figure 1* (from the VoiceXML 1.0 specification).

FIGURE 1

A Simple VoiceXML Example

```
<?xml version="1.0"?>
<vxml version="1.0">
  <form>
    <field name="drink">
      <prompt>Would you like coffee, tea, milk, or nothing?</prompt>
      <grammar src="drink.gram" type="application/x-jsgf"/>
    </field>
    <block>
      <submit next="http://www.drink.example/drink2.asp"/>
    </block>
  </form>
</vxml>
```

One of the most important concepts in VoiceXML is the *form*. Forms present information and gather input from the user. A form has *field* components that contain information for requesting a specific piece of information, such as a type of drink, as in this example. A *prompt*, or utterance spoken by the system to the user, is defined in the *prompt* element. In this example, the developer has decided that the prompt is to be spoken using the voice browser platform's built-in text-to-speech capability, so the only thing that needs to be specified is the actual text to be spoken. Recorded audio can also be specified using an *audio* element. In addition, if finer control over how the text is spoken is needed, text-to-speech markup elements can be included to perform such tasks as inserting pauses and emphasizing specific words.

Once the prompt is spoken by the system, the user responds with speech. The grammar defined in "drink.gram" is used by the speech recognizer to recognize the user's response. We do not see the grammar itself in this example; it is referred to by the relative universal resource locator (URL) "drink.gram," although it is also possible to include the grammar directly in the markup without an external reference. Once the speech has been recognized, the *block* is executed. In general, a block contains non-interactive executable code. In this case, the block specifies the next step to take in the dialog, i.e., moving to the "drink2.asp" document.

This example supports the trivial dialog in *Figure 2*.

Other Supporting Standards

Although this paper focuses on VoiceXML, other standards are also required in order to have fully interoperable speech applications. The W3C Voice Browser Working Group is currently involved in defining standards in the following areas:

1. ***Speech Recognition Grammars:*** The most significant gap in VoiceXML 1.0 was the lack of standardization of speech grammar formats. Speech grammars are the means by which developers specify what the system expects users to say at any point. With current speech recognizers, the system is essentially "deaf" to anything the user says outside of the grammar. Consequently, the specification of the grammar is crucial to the success of an application. The speech-recognition grammar standard defines two standardized formats for specifying grammars: an XML format that must be supported and an optional ABNF format that

FIGURE 2

A Sample Speech Dialog Supported by a VoiceXML Script

System: Would you like coffee, tea, milk, or nothing?

User: Milk.

System: (continues in document drink2.asp) is less verbose than the XML format. In addition, another specification is in progress that defines a more flexible type of speech-recognition grammar: the NGram format, which is not yet used by most current speech-recognition systems, but which will probably be available in the next one to two years.

2. ***Speech Synthesis:*** A standard for annotating text to be presented to a speech synthesizer with tags to direct how the text should be pronounced. For example, this standard describes tags for pauses, tags to indicate how loud the speech should be, and so on.
3. ***Natural Language Semantics:*** A standard for defining how to express the meaning of user's speech in a standard XML–based format, as well as other information about how it was spoken and when it was spoken.
4. ***Call Control*:** This standard will define features that support more advanced call control in Voice Browsers than the simple transfer and hang-up capabilities that are available now.
5. ***Reusable Dialogs:*** This standard will define how to develop reusable dialogs that accomplish specific tasks, such as getting a date or time from a user.
6. ***Speech Lexicon:*** This standard will define standards for how to represent the pronunciation of words. This lexicon could be used by text-to-speech systems to decide how to pronounce words as well as speech-recognition systems to decide how the words that it is listening for are pronounced.

The W3C Voice Browser Web site [4] maintains the documents that define the current state of these standards. As these standards mature, they will become incorporated into commercial VoiceXML browsers, providing significant improvements in interoperability.

Limitations and Cautions

While VoiceXML is dramatically reducing the complexity of developing speech applications, it is important not to lose sight of the fact that there are still many complexities remaining in this process. Some of these, such as the human factors of spoken dialogs, will be addressed, although not solved, by future versions of dialog markup languages, which will build in features that incorporate known principles of speech interface design. Other complexities, such as the issues that arise during the deployment and provisioning speech applications, scaling, and load balancing, are outside of the scope of VoiceXML or any future dialog markup language. They will be addressed primarily as the industry develops increasing experience with large-scale deployments.

As another caution, it should also be pointed out that while VoiceXML might appear to be simple to use, it is also very important when considering a commercial VoiceXML-based deployment to work with an established vendor that is familiar with the technology. The vendor must have the specialized knowledge and experience that it takes to develop telephony-based speech-recognition applications in general and VoiceXML applications in particular. As just one example, the human factors of dealing with voice dialogs are very complex, and it is easy to make human interface mistakes that can cause a project to fail. Without the specialized knowledge of speech-application experts, projects will flounder as the team members acquire the requisite expertise through trial and error.

Future Developments in Speech Applications

The future of VoiceXML and speech applications in general is very exciting. Advances will occur along multiple dimensions, simultaneously widening the scope of dialog applications, making them more natural and capable, and making them easier to design and deploy. Here are some new directions that we are likely to see:

1. ***Multimedia:*** Speech-only applications will be supplemented with integrated text and graphics as devices become more capable and as the technologies for integrating speech and graphics become more mature. Graphics is a better output device than speech for many types of information, such as maps and diagrams, and the only choice for information that must be conveyed in the form of images.
2. ***More Natural Dialogs:*** Current VoiceXML applications, although effective, generally tend to have a machine-driven feel, despite the fact that VoiceXML itself contains features that support mixtures of machine and user-driven dialogs. Major roadblocks to more natural dialogs include (1) current speech-recognition technology performance degrades as the number of things that the recognizer has to listen for increases, (2) difficulty to design natural dialogs with current tools, and (3) difficulty to design and debug the complex speech-recognition grammars that support natural dialogs. All of these factors will change as speech-recognition technology and dialog design tools improve. Better speech-recognition technology also will allow systems to take advantage of much more flexible and easier to design language models, such as NGram models, which will eliminate the problem of debugging speech grammars.
3. ***Better Speech Recognition and Text-to-Speech Technology:*** Speech-recognition errors will never go away entirely; even humans sometimes have trouble recognizing the speech of other humans, and there is no reason to think that machines will be able to do better than people. Sometimes there just is not enough information in the speech to tell what someone is saying. Humans have developed many strategies for effectively communicating with each other under poor conditions, and many of these strategies are beginning to be exploited by sophisticated computer dialog designers. However, there is still a great deal of room for improving the baseline accuracy of speech recognizers. Think about a conversation with a person who is hard to understand for whatever reason. For example, the person you are talking to might be speaking a foreign language that you don't know very well. Even though you can usually manage to communicate, the conversation is much easier and more pleasant when speech recognition is not an issue. Text-to-speech technology is also making tremendous improvements and sounds much better than it did only a few years ago. Nevertheless, it still sounds artificial, and serious commercial deployments use recorded natural voices as much as possible.

4. ***More Devices:*** VoiceXML browsers are primarily based on telephone platforms today, except for browsers used by a few desktop testing tools. However, a very natural next step would be to use VoiceXML browsers on PDAs and desktop computers, especially as multimedia capabilities for VoiceXML develop.
5. ***Better Tools***: Just as it is possible to design a graphical Web site with only a text editor, it is also possible to develop a full VoiceXML application. This is part of the appeal of VoiceXML. No specialized tools are needed get started. But manually managing a complex commercial VoiceXML site that not only may include VoiceXML pages, server-side code, sound files, and speech grammars, but must also interoperate with existing integration code that supports a graphical Web site, using only simple tools, is extremely difficult. Commercial design tools that reduce the complexity of the VoiceXML development process are starting to be developed. In addition to the choice of a VoiceXML platform itself, a VoiceXML development project should also give careful thought to the available tools and make a considered decision on which development tools to use based on well-understood requirements. Tools range in capability from simple XML editors, through more intelligent VoiceXML specific editors, to abstract dialog design tools that developers can use without even being aware that they are developing VoiceXML applications. Using these tools will dramatically accelerate the adoption of VoiceXML.

Conclusions

VoiceXML is an exciting technology that is not only starting to open up the Internet to telephone access, but is also dramatically reducing the development time for telephony-based speech applications. Even though the first specification was released only two years ago, VoiceXML is rapidly becoming a major factor in the speech/telephony marketplace that deserves serious consideration.

References

1. V-Commerce: The Next Wave: www.finfacts.com/lock/ feb2/11/VR.htm
2. VoiceXML Forum: www.voicexmlforum.com
3. World Wide Web Consortium: www.w3.org
4. World Wide Web Consortium Voice Browser Group: www.w3.org/Voice
5. XML: www.w3.org/XML
6. Christensen, Clayton, *The Innovator's Dilemma.* (Boston: Harvard Business School Press, 1997).

Meeting IP Service Provisioning Challenges

Anand Desai
Vice President, Marketing
DSET Corporation

IP Service Activation and Provisioning: A Journey with Great Potential, Major Challenges

When Portuguese navigator Ferdinand Magellan set sail in 1519 on a quest to find a western route to the Spice Islands, few could have known that his expedition would establish the world's first true communications network.

Magellan's ships crossed the Pacific, became the first fleet to circumnavigate the globe, and ultimately proved the roundness of the Earth.

Flash forward nearly 500 years, and today's "explorers" often take the form of telecom engineers, also on a mission—focused on delivering voice, data, and other information around the very same world as Magellan. For both journeys, provisioning was key.

"Provisioning" is a term that has long been used in the telecommunications industry to designate all of the network operations that must be completed to achieve one objective: turning on services for a telecom provider's customers. At one time, the sole focus of provisioning was to turn on POTS, or plain old telephone service. Today, however, there are numerous types of providers offering many different voice and data services, with each service holding the promise of substantial financial return in addition to presenting its own set of technical challenges.

At the start of the 21st century, Internet protocol (IP)–based services, such as virtual private networks (VPNs), are increasingly attractive to service providers and their customers alike. For providers, these services present a significant growth area in a hotly competitive marketplace. As with any telecom service, efficient provisioning can be the key to success in this market, as it can reduce costs substantially while enhancing revenues. But in the rush to surge ahead of competition, service providers must not be shackled by poorly designed provisioning solutions that limit flexibility and growth. This can only result in the failure to meet service demands, customer expectations, and revenue targets.

Multiple Services in a Multivendor Environment

Success demands flexible provisioning solutions capable of supporting multiple services in a multivendor operations environment. Such solutions must also be highly scalable, robust, and extensible to meet new market needs. Building a monolithic piece of software to address the requirements imposed by one service or a single equipment vendor severely limits a provider's ability to keep pace with the evolution of the market. However, a product whose design is based on flexibility, modular components, and service definitions will enable a provider to maximize return on investment (ROI) despite the most fluid technical and commercial conditions. A service-, device-, and vendor-neutral solution is a requirement for emerging networks.

Pitfalls of Manual Provisioning

While simple in concept, provisioning is often complex in execution, especially when it has to be done manually. Essentially, provisioning is initiated by a customer's service order to a communications provider and culminates when the customer is able to use the service purchased. During this process, the service is enabled for use on one or more devices in the network. Turning on a VPN service, for example, may entail such complexities as sending commands to several network routers.

In today's telecom environment, manual provisioning is slow and error-prone. Even provisioning a small VPN can be time-consuming. Manual provisioning also requires constant vigilance by highly skilled operations personnel attuned to myriad difficulties that can arise from changes in vendor products and industry standards. For example, the inability of digital subscriber line (DSL) to grow beyond a certain market size can be attributed to lack of cost-effective automated provisioning.

In addition, there are the problems caused by partially failed service activations. Partial failures can compromise the integrity of the service infrastructure to the point where costly manual intervention is needed. Manual rollback of a failed multisite VPN service order can be a service provider's nightmare.

A Clear Imperative

Clearly, automated IP provisioning is imperative for the success of any communications provider competing in this market segment, as well as for large enterprises seeking to manage their own networks as efficiently as possible. However, building a provisioning system that can deliver all of the capabilities critical in a multi-technology, multivendor environment is a daunting task.

To begin with, providers typically offer a variety of services supported on different types of servers, routers, firewalls, and switches. For each type of device, there may be multiple vendors. And to further complicate matters, change is always around the corner, which may be a version upgrade on a router, a new router model, a change in the way a service is deployed, or an entirely new service.

The provisioning systems currently on the market have been designed to be single-service or single-vendor solutions and lack essential flexibility and extensibility. Service providers are thus forced to buy different systems for different services and vendors. Resources cannot be shared by multiple services. Having multiple systems makes routine operations as well as maintenance far more expensive and complex than they would be if one system could be deployed to automate provisioning—regardless of the service or vendor.

Provisioning Services, Not Devices

A critical deficiency in conventional provisioning systems is that they are based on a myopic design philosophy that views the provisioning process as "device provisioning." A more elegant and effective approach is to view provisioning as enabling a service for the end user. This perspective reflects the key direction in which the market is moving. Subscribers want service-level agreement (SLA) compliance that essentially implies quality of service (QoS) guarantees. It also implies that the essence of "customer satisfaction" is developing a service model that can be rapidly deployed and that accurately represents a subscriber's expectations for that service.

A provisioning system should also be the repository for provisioned data. This is a complex requirement, since such a repository has to have comprehensive information about the service provider's own equipment, organization, and internal users as well as each subscriber's organizational structure, devices, and authorized users. Additionally, this information has to be accessible for queries from other operations support system (OSS) components that have their own fixed views of the service infrastructure and state. For example, some OSSs will issue queries based on subscribers, while others will generate queries based on devices.

Building a Better Provisioning Solution

Effectively addressing these new types of design issues requires a component-based, modular "plug-and-play" approach to building a service-provisioning system. Developing a system that is flexible and extensible also requires a clear separation of service applications from a provisioning core and device models.

A service application consists of service and policy definitions. Service definitions enable activation of the service by the provisioning core, and policy definitions include the types of policies activated for the service. This will, for example, include QoS policy definitions where policy rules relating to traffic metering/monitoring will be part of the policy definition for a VPN service.

A provisioning system should allow the dynamic introduction of new service definitions, policy definitions, and device models. This will enable the service provider to travel a very fast time-to-market path.

Another design consideration is how well a service-provisioning solution fits into the overall OSS environment, where the provisioning system will have to interact smoothly with a variety of back-end systems. For example, well-designed provisioning software must be capable of interfacing with performance-management systems so that SLA profiles can be populated with service-configuration information, such as the set of interfaces required for a VPN service. It is also desirable that the provisioning system be able to communicate with an inventory system to import information about available network devices. Similarly, a fault-management system may have to query the provisioning system to determine the service configuration for a particular device.

A Universal Provisioning Solution

A universal provisioning solution, carefully designed, can provide all of the capabilities indispensable for fast, cost-effective IP provisioning. A highly scalable, standards-based, robust, and extensible architecture is needed that enables rapid, network-wide, policy-based, and error-free flow-through provisioning of IP–based VPNs for multivendor networks (see *Figure 1*). Integration into an OSS environment by means of open interfaces into order-entry, billing, and management systems is also needed to further automate the entire management process.

Furthermore, a rich modeling of service applications, and drivers for communicating with all of the network elements that have to be provisioned and managed is needed. A platform-based approach that offers all the features of a transaction processing system, such as reliability, scalability, availability, and rollback upon failure can certainly address the service-activation needs of millions of subscribers of many blends of services.

Key Features of a Carrier-Class Service-Activation Solution

Customizability and Extensibility

A system that allows maximum customization to support new or modified products, services, and networks without any major change is needed. Service applications based on a rich modeling of services and service models should be easily added or edited as required.

Scalability

Service providers demand scalability to address the volume of transactions as business grows. Bulk provisioning, and

FIGURE 1

IP Service Provisioning with UPX

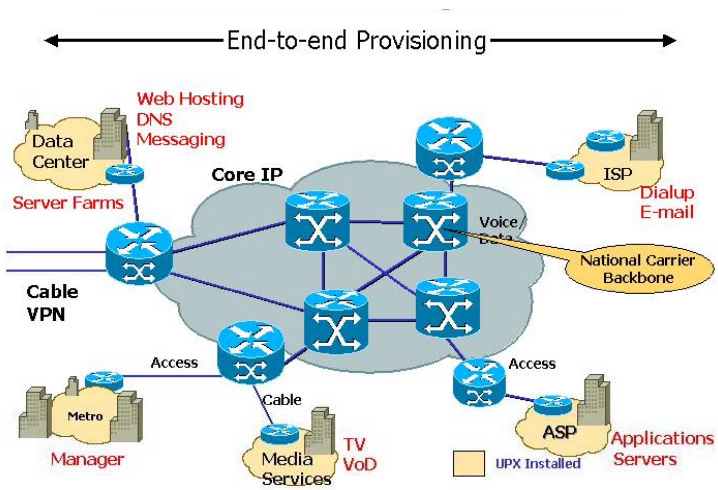

the ability to handle millions of subscribers in real time, is the challenge that needs to be met.

Service Bundling
Bundling of existing service applications can enable the seamless introduction of new services. For example, an IPSec (IP security) VPN service may be bundled with a QoS application to significantly cut the time to market for an IPSec QoS VPN offering. These combinations can help create value-added services needed for real-time applications used in business, entertainment, and gaming.

Policy Support
Users should be able to create policy templates, such as those for security, that facilitate the creation of policy rules to be applied during provisioning. The classification of policies allows the activation manager to easily create, customize, and apply policy rules.

OSS Integration
An open architecture that supports interfaces based on extensible markup language (XML), common object request broker architecture (CORBA), and Java messaging services (JMS) to maximize interoperability with other OSSs, including order-entry, billing, inventory, and fault-management systems, can blend into an existing OSS ecosystem. Open and flexible interfaces can make the typically difficult OSS integration process a much simpler exercise.

Vendor-Neutral Provisioning
A modular and vendor-neutral design can provide support for provisioning a wide variety of devices.

Transaction Support
The network is a financial asset that needs to be utilized appropriately. Each provisioning request should be treated as an important transaction. Besides, a number of network requests may be generated from a single provisioning request. The failure of any single network request is a failure of the entire transaction. Upon a failure, therefore, automated ways to initiate rollback requests are needed to ensure that the state of all devices affected reverts to that existing before the start of the transaction.

Security
It is important that each operator is authenticated and can perform only authorized functions. In addition, security between different components should be configured to enable encryption of messages and two-way authentication through certificates. No message should be viewed in transit, and connections need to be authenticated.

Customer Self-Provisioning
A service provider's clients should be able to provision and manage their own services—a capability that can also reduce a provider's operational costs. Client self-provisioning should be securely configured so that individual clients are authorized to access and exercise control only over certain network systems and information.

Standards-Based
To help maximize utility and flexibility for service providers, the architecture should be based on industry standards such as CORBA, XML, J2EE, lightweight directory access protocol (LDAP), and simple network management protocol (SNMP). The service-application standards reflected in this architecture are IPSec, Internet key exchange (IKE), multiprotocol label switching (MPLS), differentiated services (DiffServ), and remote authentication dial-in user service (RADIUS).

Conclusion

Automated service-activation and provisioning solutions that provide carrier-grade features with the flexibility and extensibility are critical for meeting the fast-changing commercial and technical demands in the communications industry. The need has never been greater—particularly as new services and devices are continuously added to the mix.

An open architecture can further enable seamless interaction between the different components of a service provider's OSS environment. Highly scalable and robust, these solutions can minimize the money, time, and manual effort that providers need today to customize services, to bundle existing offerings to create new ones, to quickly introduce completely new services, to deploy equipment from new vendors, or to upgrade equipment supplied by current vendors.

Automated provisioning is a quest worthy of Magellan – a voyage into still uncharted seas with rich possibilities on the other side of the globe.

Cities of the Future: The Role of Smart and Sustainable Communities

John M. Eger
Van Deerlin Professor of Communications and Public Policy
President of the California Institute for Smart Communities
San Diego State University

Cyberplace and Cyberspace

According to author Charles Handy, we live in an age of paradox. The more high-tech our world, the more high-touch we are becoming. The more global, the more intensely local our focus needs to be. The more competitive our markets, the more cooperation is a critical element in developing our business strategies.

One of the more interesting paradoxes of our time is that the more we live and work in cyberspace, the more important real place becomes. While this notion runs counter to much of today's popular literature, we are already seeing the knowledge worker and the high-tech knowledge-sensitive industries migrating to highly livable communities—communities with mountains or lakes, open spaces, clean air and water, and, as in the case of Portland, Oregon and other communities where they have established urban growth boundaries, less reliance on the automobile as the primary mode of transportation.

This growing concern with urban sprawl, coupled with the nostalgic yearning that the new urbanism movement represents, are evidence of sweeping changes in public attitude toward physical space. As the Internet revolution moves into full bloom, however, there is every reason to believe it will have a dramatic impact on the architecture and landscape of communities throughout the world.

No technology in human history is having, or is likely to have, such tremendous influence on life and work and play, and, in the transforming process, on our physical space. While a "smart community"—a community that makes a conscious decision to aggressively deploy technology as a catalyst to solving its social and business needs—will undoubtedly focus on building its high-speed broadband infrastructures, the real opportunity is in rebuilding and renewing a sense of place, and in the process a sense of civic pride.

Athens, the place where civilization was born and where the city-state form of governance first began, remains a symbol of the dynamic potential of cities to create and provide the linkages among culture, commerce, and civic pride that are so important to the wealth and well-being of a community. Over the years, however, cities have been both cursed and blessed as they have been compelled to adjust to the underlying changes taking place in our movement to a global economy and society. Many cities have already died; others are in fiscal and societal decay.

The concept of cities as engines of civilization remains deeply embedded in our collective psyche. As cities of the past were built along railroads, waterways, and interstate highways, cities of the future will be built along information highways—broadband communications links to homes, schools and offices, hospitals, and cultural centers, and through the World Wide Web (WWW) to millions of other locations all over the world.

As past is prologue, surely some cities will become the ghost towns of the 21st century information age. By far, however, cities will succeed and survive in this next transition to a knowledge-based, global information economy and society. Indeed, cities of the future—the smart and sustainable communities built for the digital age—will play a central role in the rebirth of civilization in the 21st century.

The Power and Pervasive Influence of Technology

In less than a decade, the great global network of computer networks called the Internet has blossomed from an arcane tool used by academics and government researchers into a worldwide mass communications medium, now poised to become the leading carrier of all communications and financial transactions affecting life and work in the 21st century.

The growth of the Internet's World Wide Web has been even more spectacular. With 30 million-plus users worldwide, growing at 15 percent per month, it is being integrated into the marketing, information, and communications strategies of nearly every major corporation, educational institution, political and charitable organization, community, and government agency in the world. No previous advance—not the telephone, the television, cable television, the videocassette recorder (VCR), the facsimile machine, nor the cellular telephone—has penetrated public consciousness and secured such widespread public adoption this quickly.

In recent years, it has become fashionable to refer to the domain in which Internet-based communications occur as cyberspace—an abstract communication space that exists both everywhere and nowhere.

But until flesh-and-blood human beings can be digitized into electronic pulses in the same way in which computer scientists have transformed data and images, the denizens of cyberspace will have to live IRL (in real life), in some sort of real, physical space—a physical environment that will continue to dominate our existence and our future in the same way that our homes, neighborhoods, and communities do today. Nonetheless, information technology (IT) is a force that will reshape our world as never imagined.

The Rise of Smart Communities

Already, communities and nations around the globe—often without being consciously aware of it—are starting to sketch out the first drafts of the cyberplaces of the 21st century. Singapore has launched its IT–2000 initiative, also known as the Intelligent Island Plan. Japan is building an electronic future called Technopolis, or Teletopia. France, as early as 1976, initiated a plan called Telematique, an aggressive effort to place personal computers (PCs) on every desktop and in every home in the country. In the United States several years ago, the former Clinton Administration pursued a vigorous National Information Initiative, or NII, one of whose early goals was to link every school and every school-age child to the Internet by 2000.

Many communities—Stockholm, Seattle, and Sacramento, for example—have constructed large-scale public-access networks that residents can use to obtain information about government activities, community events, and critical social services such as disaster preparedness, child-abuse prevention, and literacy education. The tiny university town of Blacksburg, Virginia has transformed itself into an electronic village in which the majority of the town's businesses and residents are connected to the local data network. And counties such as San Diego, as a result of its City of the Future project, are building even more sophisticated electronic infrastructures that, one day soon, will allow a wide variety of local government, business, and institutional transactions.

Recognizing that electronic networks such as these will play an increasingly important role in a municipality's economic competitiveness, the State of California four years ago launched a statewide Smart Communities program, which has been managed since its inception by the International Center for Communications at San Diego State University. More recently, a World Foundation was established to help other communities around the world with their struggle to get on the global information highway.

The Foundation defines a smart community as a geographic area ranging in size from a neighborhood to a multi-county region whose residents, organizations, and governing institutions are using IT to transform their region in significant, even fundamental ways. A fundamental premise is that smart communities are not, at their core, exercises in the deployment and use of technology, but in the promotion of economic development, job growth, and an increased quality of life. In other words, technological propagation of smart communities is not an end in itself, but only a means to reinventing cities for a new global information economy and society with clear and compelling community benefits.

Technology, Culture, and Place

One of the main reasons that information networks can have such a profound transformative effect on people, businesses, and communities is that every other major technology advance that has shrunk space and time has also remade society in fundamental and important ways. Transportation, over the millennia for example, has done more than perhaps any other technological advance to bring the world's people closer together.

But telecommunications developments, including telephones and their more modern kin, accentuate the trends inaugurated by transportation advances in three slightly different, but very important, ways.

First, by allowing for rapid communication between distant sites, they make it possible for business and social relationships to flourish over long distances, permitting workers and investment capital to migrate to the most desirable locations and those with the highest economic return.

Second, they extend the reach of these economic, social, and other relationships far beyond national borders, creating what is truly a global economy. And third, and perhaps most significantly, they make possible for the first time the nearly instantaneous transmission of information, collapsing both space and time in a way that no other previous technological advance has done.

The Internet, the World Wide Web, and their successors are likely to produce consequences that are as great or greater than anything we have seen so far—and that are apt to be equally unexpected. If we are to maximize the positive contributions of these new technologies while minimizing their negative ones, we must begin to appreciate now how these technologies are likely to affect our people, our culture, and our perceptions of place in the years to come.

The Technical Architecture of the Smart Community

There are a few general trends worth noting. The first is the growing ubiquity of telecommunications networks. Because it is based largely on the existing telephone systems, the Internet today spans the globe, with its tentacles reaching into more than 130 countries and connecting, in one form or another, an estimated 400 million to 450 million people. This expansion shows no signs of abating. Indeed, as the Internet migrates from its almost purely copper-based telephone platform to cable, satellite, and digital cellular systems, the methods of connecting to the Internet will proliferate, access costs will decline, and the number of users will skyrocket.

The second general trend in the development of the Internet is the rapid expansion in bandwidth. In its original incarnation (which lasted for more than two decades), the Internet was primarily a low-volume text-based medium, and so required little transmission capacity.

The emergence of the World Wide Web, with its heavy use of graphics, photographs, and animation, changed this equation dramatically, and even top-of-the-line modem technologies—28.8 kilobits per second (kbps) and even the 56 kbps modems—quickly proved inadequate to the task of transporting these billions of bits of graphical information, causing many parts of the Internet to react like a two-lane freeway suddenly jammed with a hundred- or thousand-fold increase in the number of vehicles.

The third and perhaps most important trend in the development of the Internet is the proliferation of access points. Until recently, logging on to the Internet has required a fairly sophisticated computer, costing in the neighborhood of $1,000. While this is one-half of what it was just two years ago—this expense has still priced the Internet out of the range of a large share of low- and middle-income families in the United States, not to mention the vast majority of the rest of the world's population.

This high cost of access has combined with the relative inconvenience of using a computer—sitting before a computer, unlike a television, is hardly the most relaxing experience—to restrict the Internet largely to the technologically oriented, well-to-do minority. This is one of the main reasons why many communities have undertaken aggressive public access initiatives to install computers and kiosks at community centers, public libraries, and other public sites to make it possible for people who do not own a computer to use the Internet.

But this situation also is changing. Already, several companies, including Sony and Phillips, have introduced devices that allow people to log on to and browse the Internet directly from their televisions, and the number of such devices is likely to multiply during the next two years, particularly as cable-television companies become more involved in the Internet-access business. Similarly, other companies are beginning to distribute videoconferencing equipment that will allow people to make videophone calls over the Internet to and from their televisions.

As a result of developments such as these, we are quickly reaching the point where the world will be interconnected by a next-generation Internet that allows for instantaneous transmission of text, photographs, and broadcast-quality audio, video, and virtual reality, and not necessarily to expensive computers nor any other new technological device, but rather to the ordinary televisions that are now in place in hundreds of millions of living rooms worldwide.

The Changing Geopolitical Context

These technological changes are taking place at the same time that the world's geopolitical landscape is being radically redefined. No longer dependent upon national governments for policy ideas and information, and no longer content to be bound by the one-size-fits-all pronouncements of national legislators, local leaders are taking social and economic matters into their own hands, pursuing policies that will promote job creation, economic growth, and an improved quality of life within their region, regardless of the policies enacted at the national level.

This reverse flow of sovereignty, by which local governments are assuming more responsibility than ever before for their residents' well being, has come about at a time when information and markets of all types are becoming increasingly globalized. News, currency, and economic and political intelligence—not to mention products and services—can no longer be contained within national borders. Increasingly, they flow, often instantaneously, to all corners of the globe, making it difficult or even impossible for national governments to influence political situations or economic conditions over which, not long ago, they held unquestioned control. The result is a geopolitical paradox in which the nation-state, too large and distant to solve the problems of localities, has become too small to solve the borderless problems of the world.

Locally based companies that once competed with firms only in their own area code, for instance, now must battle companies throughout the world for their customers' loyalty and dollars. Local governments that once had to compete for high-value residents against only nearby municipalities and the amenities they could muster, now must struggle to attract such residents in a world where a growing number of people can live nearly anywhere they want and still have access to the same jobs, the same income, and the same products and services to which they have grown accustomed.

To meet these challenges, many far-sighted localities have begun to transform themselves from fractured, often highly contentious regions in which a thousand interests compete for larger shares of a shrinking pie into something more akin to the city states of old than to the archetypal municipalities of modern-day political-science texts.

Those that are succeeding, such as Smart Valley and San Diego in the United States; Stockholm, Hong Kong, and Infoville in Spain; or Malaysia's Multimedia Corridor possess a number of common features. One characteristic is collaboration among different functional sectors (government, business, academic, nonprofit organizations, and others) and among different jurisdictions within a given geographical region. These collaboratories are fast becoming the new model for successful urban organization in the global age and the only local political arrangement likely to make it possible for besieged municipalities to survive the increasingly intense global competition that lies ahead.

This subtle point is often lost in discussions about building smart communities and even in the implementation of many of the smart community projects themselves. But it could not be more important. Indeed, proponents argue that the more time people spend in cyberspace, the more important real place becomes, and the more civic involvement and the real values of communities—places where common dreams and visions become reality—become apparent to success and survival in this new age.

This new competitive and community spirit, however, will not come about automatically. Communities must develop a coherent and compelling vision that makes it clear how the new information networks are going to promote job growth, economic development, and improved quality of life within the community, and they must communicate that vision broadly. This is the key element that is missing from so many smart community plans today, and yet it is the most essential: for, unless a community knows precisely where it is headed and how it hopes to get there, it is unlikely to reach its destination—to its detriment and to all who are stakeholders in this new but uncertain future.

Toward a Philosophy of Permanence

Fortunately, a new breed of architects, planners, and developers is beginning to pencil in that new vision of cities in the information age. It is a bold vision that deals with the crises of growth and the current development sprawl, while returning to a cherished icon—that of a "compact, close-knit community," according to Peter Katz, author of *The New Urbanism: Toward an Architecture of Community*.

The prospect of a new century, says Katz, raises serious concerns. Former U.S. Vice President Al Gore believes we are on a collision course between our worldwide civilization and the ecological system of the earth.

Many policy wonks agree that the urgency of our dilemma has reached an acute stage. Thus, as we examine our current policies of land development and urban planning, new nonlinear solutions are imperative. The thing that we must remember, urges Katz, is that all of the strategies must be examined, tested, and tested again in relation to prevailing developmental models. Only then can we determine if a new urbanism can indeed be shown to deliver a higher, more sustainable quality of life to a majority of this nation's citizens.

One of the more interesting and exciting aspects of the new urbanism movement is that the next paradigm could well be much more than the return to the close-knit community of the small town with its village greens and mixed-use zoning. It could be a spiritual return to the kind of community enjoyed by the earliest cities.

Tessie Naranjo of the Santa Clara Pueblo in New Mexico defines community as " the human dwelling place." It is where the people meet the needs of survival and where they weave their webs of connections. Native communities are about connections because relationships form the whole. Each individual becomes part of the whole community, which includes not just the human population, but also the hills, mountains, rocks, trees, and clouds.

Until recently, advances in telecommunications and transportation have contributed to our disconnected-ness, rather than cemented us as a people—they have atomized our sense of community rather than provided us a sense of place. Yet without a cultural center, a shared history, or a commitment to neutral goals and visions, there is little to cement communities together.

Chief Sealth, for whom the city of Seattle is named, cautioned: "This we do know: the earth does not belong to man, man belongs to the earth. All things are connected like the blood that unites us all. Man did not weave the web of life; he is merely a strand in it. Whatever he does to the web, he does to himself."

As the World Wide Web becomes part of the web of life, perhaps mankind's technology will ultimately enhance and secure our connectedness to the physical world, preserving and protecting it for future generations. If successful, the smart and sustainable community will dramatically reverse an adverse trend precipitated by the invention of the cotton gin and the industrial revolution that followed, by the automobile and the fifty years of untamed growth and land development, and, worse, by the advance of a rootless culture without a sense of place and will help to lead us out of the spiritual and physical wasteland that we have created.

Conclusion

Many cities—some deliberately, others almost unconsciously—are already benefiting from the "greening" of their cities, the reinvestment in downtowns, and the high-tech information infrastructures of their urban planning.

Because nobody knows for sure what the city of the future will look like—and we are almost never writing on a clean slate, that is, launching a brand new city out of virgin territory—it makes good sense to paint our vision of the future with a broad brush but to be careful to preserve and protect our communities unique character, history, and culture.

For example, we must recognize the value of our existing investments—often downtown and in historic areas of the community wherein lies our "cities' DNA," according to Kevin Starr, historian, author, and California State Librarian. In those vital centers, we must steer both our public and private investments and encourage public/private partnerships to renew and redevelop for an information age.

Downtowns seem to be a particularly attractive haven for redevelopment and attracting the multimedia industries. Denver, with its Beaux Arts Historic Court, is halfway through a decade-long downtown investment of one billion dollars in three sports stadiums, an aquarium, and an urban shopping complex. According to a report by *New York Times* reporter James Brooks, "Denver, following retail and entertainment development downtown, is now earmarking one billion dollars in residential construction. A total of 1,334 apartments is being built or converted downtown, more than the total of the last four years."

A 1999 study by the Brookings Institution, indicated this is part of a national trend, at least in the United States. With mixed-use zoning, libraries, shops, and live/work lofts all within walking distance of one another, these communities are fast becoming the most livable communities in the age of the Internet.

Another aspect of a new agenda entails ending the "fiscalization" of land use—a new term coined to describe the extent to which tax-based concerns drive land-use decisions. Why? All too often, local officials fear that "if pushed too hard for alternatives to sprawl by requesting improvements in a development project, the developer or company involved will sue them or move on to the neighboring jurisdiction," according to Constance Beaumont of The National Trust for Historic Preservation. Beaumont explains in the publication *Smart States, Better Communities* that "given their heavy reliance on property taxes for local revenue, most municipal governments find it extremely difficult to say no to any new development, even if it means losing their identity and even their way of life."

Fortunately, the fundamental shift to a knowledge-based economy and society is fast becoming a reality. This movement, in which wealth creation hinges on our ability to create knowledge-based information products, ushers in a new era in which sophisticated electronic infrastructures allow a wide variety of government, business, and institutional transactions to take place electronically and obviates the need for travel to and from the government agencies and other places of business.

Importantly, too, such infrastructures are what are needed to attract the knowledge industries of the new information-based economy. What we are also beginning to understand is that the smart and sustainable, or livable, communities also attract the knowledge workers themselves, who are migrating to places such as Portland with its urban growth boundary and Seattle with its lakes and trees, and precious open spaces.

The agenda for retooling for the information age is enormous. Education, health care, travel and tourism, art and culture, sports, business, and government all need to be redefined in this new age, and all are vital to creating the "whole-brain" community—responsive to and reflective of the kinds of people that we are and must be to succeed in the knowledge-based global economy challenging us on every front.

Communities that want the economic and social advantages that come with being a smart community must work together to create a new vision of that future for themselves, must use the political process to engage the whole of their community in that effort, and must commit to transforming work and play, government and business for this new age. In a way, becoming a smart community is really not so much about deploying technology to reinvent the community as it is about engaging the body politic to reinvent governing in the digital age.

We have the tools of this new age—computers, telecommunications, and information appliances of all kinds. We have the software, too. Indeed, because of our First Amendment and free enterprise culture, we produce more books, movies, and software programs than any other nation in the world.

But most of all, we need people—each other—and places where culture and commerce, and civic pride are joined—where we are refreshed, energized, and challenged to achieve the best we are capable of. Cities do that; they provide the context for our lives and the fabric of our existence—our actions and our enterprises.

Information-age cities of the future, such as Athens of an earlier age, will be truly smart communities—sustainable, healthy, culturally strong, diverse, and exciting places to live and work, and play.

References

1. *Building Smart Communities: How California's Communities Can Thrive in the Digital Age* (San Diego State University, International Center for Communications, January 1997).
2. *Building Smart Communities: Smart Communities Implementation Guide* (San Diego State University, International Center for Communications, January 1997).
3. "Cities, Smart Growth and the Urban Future," *San Diego Union-Tribune*, February 13, 2000, Insight.
4. *City of the Future: Report of the Mayor's Advisory Committee on the City of the Future* (San Diego State University, International Center for Communications, 1994).
5. "Cyberspace and Cyberplace," *San Diego Union-Tribune*, October 26, 1997, Opinion.
6. *San Diego in the Global Village: Telecommunications Infrastructures Study for the City and County of San Diego* (San Diego State University, International Center for Communications, 1992).
7. "Wiring the Future," *San Diego Union-Tribune*, July 2, 1995, Opinion.

Heading toward Next-Generation Public Telephony

Enabling Universal Telephony through New Infrastructure Equipment and Signaling Protocols

Didier Heen
Vice President, Marketing and Partnerships
CIRPACK

Operator telephony infrastructures are moving to a new model based on open technologies. This revolution is comparable to the changes experienced by the information technology (IT) industry a decade ago, when mainframes began to be replaced by client-server equipment. At the same time, we are seeing the new broadband local loops gaining market-share. Digital subscriber line (DSL), cable, local-area network (LAN), wireless, etc., are more suitable for service providers in order to offer both high-speed Internet access and voice communications at a low cost.

However, neither telecom operators nor their customers will switch to "next-generation" protocols and infrastructures overnight, but rather over several years, if not decades. Due to massive investment done over the years in old-generation equipment, service providers are asking for technologies that enable smooth migration of both legacy infrastructures and services to the open packet-telephony model. This is how service providers will save on costs while offering more value to customers and increasing profitability.

1. From Circuit to Packet-Based Infrastructures

Over the past 30 years, telecom operators around the world have heavily invested in time division multiplexing (TDM) equipment, building hierarchical, circuit-based infrastructures. In this legacy model, the local-exchange switch provides three main functions in one box:

1. *Subscriber Connectivity:* Analog, integrated services digital network (ISDN), and V5 LE for concentration
2. *Subscriber Services:* Providing dial tone and handling class services such as call waiting, call transfer, command line interface (CLI), etc.
3. *Core Network Connectivity and Services:* E1, signaling system 7 (SS7), and transit functions

Public TDM switches were developed in the early-1970s for state-owned Post Telephone and telegraph Administrations (PTTs) so as to provide high connectivity but basic phone services. These machines are very robust and provide high-quality voice services. However, TDM technologies are now showing their age. They are based on proprietary hardware and software that are expensive to maintain and upgrade. Introducing new services is a long, difficult, and costly process. In addition, TDM is not compatible with the new broadband local loops such as DSL, cable, wireless, or LAN, which are all based on packet technologies.

Hierarchical TDM architectures are also reaching their limits. Since each switch has to handle all the complexity of a carrier's infrastructure, more subscribers lead to more central processing units (CPUs) for services and more interfaces for wires. Equipment costs do not scale down over time, but rather continue to increase when taking on new customers. This has become very critical with the increasing use of dial-up Internet access, and this is not to mention the complexity of transit network sizing and the ineffective use of bandwidth in both the local loop and the transit networks.

Next-generation networks (NGNs) have been developed to overcome these limitations. The main features and benefits of NGNs are as follows:

- Voice and data are carried over the same transport and access networks, as packets of data rather than circuits. This allows for dynamic bandwidth allocation between voice and data services and makes much better use of resources. The transmission network is easier to scale and to maintain, and services can be provisioned more quickly.
- Intelligence needed to provide voice services is centralized rather than disseminated, reducing investments for edge devices. This accelerates network deployment to reach new customers and provides greater flexibility to upgrade or introduce new services, whatever the local loop (DSL, fixed wireless, cable, LAN, etc).
- Switching tasks are distributed rather than centralized. Voice and data are carried directly to their destination

without affecting the processing power of other equipment or loading transmission networks as in the hierarchical public switched telephone network (PSTN).

- Network equipment use open technologies and standard protocols, rather than being tied to proprietary designs. Equipment interoperates with other NGN devices, making service providers independent of a given vendor or technology. They benefit from a large choice of vendors, greatly accelerating the development timeframe and making new features easier to develop and deploy.

2. New Concepts for Voice Infrastructures

NGNs have led to the creation of new concepts for voice infrastructures, where the main functions are split into different boxes (see *Figure 1*).

The Transport Layer
The main roles of the packet phones and phone concentrator (gateways) are as follows:

- To transform voice into packets
- To terminate edge signaling
- To provide additional features, such as echo cancellation and voice compression

The Control Layer
The main roles of the device called the media gateway controller (MGC), or softswitch or gatekeeper, are as follows:

- To control the call process
- To control signaling and edge devices
- To provide subscriber and transit services
- To interface various devices with each other so as to provide advanced services

The Application Layer
The main roles of the application servers are as follows:

- To provide advanced services
- To manage service activation and mediation
- To interface with the operations support system (OSS)

The softswitch controls distributed gateways via an Internet protocol (IP) command network independent of the transport network, which is using TDM, asynchronous transfer mode (ATM), or IP transport protocols (see *Figure 2*). While the softswitch handles the entire call process and delivers basic services to subscribers (dial tone, call waiting, CLI, etc.), it needs access to external servers in order to handle advanced voice services (Centrex, virtual private network [VPN], universal messaging, phone conferencing, prepaid services, etc.).

By offering a total network layer abstraction to the control layer, this architecture opens the door to universal telephony. Operators now have the competitive advantage of being able to connect any local loop, at any time, when new opportunities arise and to provide the same services to all subscribers. By using application servers, operators can easily introduce innovative services, developed in-house or via any third-party vendor complying with industry standards.

3. Signaling as the Key to Interoperability

Signaling is the key to making universal telephony a reality. For service providers to connect any type of local loops together with any type of core network, softswitches and gateways must implement, understand, convert, or relay multiple TDMs and packet-based protocols. Building an NGN, while preserving old-generation equipment and services and maintaining interoperability with the networks of other operators, is giving rise to three different signaling challenges:

1. Transparently carrying ISUP/SS7 and ISDN protocols over IP to maintain existing services
2. Managing pure end-to-end IP telephony signaling in order to introduce innovative services and devices
3. Enabling calling between circuit (TDM) and packet (NGN) telephony devices

TDM Signaling Transparency over a Packet Network
SIGTRAN
Signaling transport (SIGTRAN) has been developed by the Internet Engineering Task Force (IETF), the international

FIGURE 1

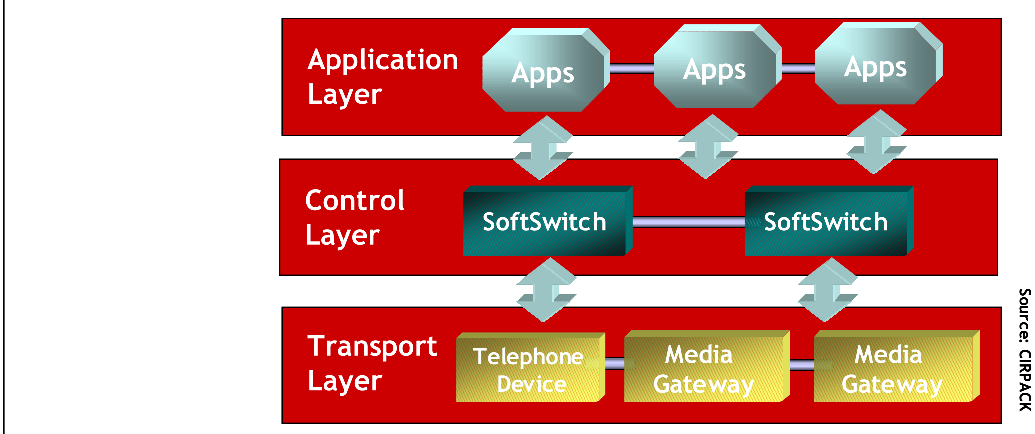

FIGURE 2

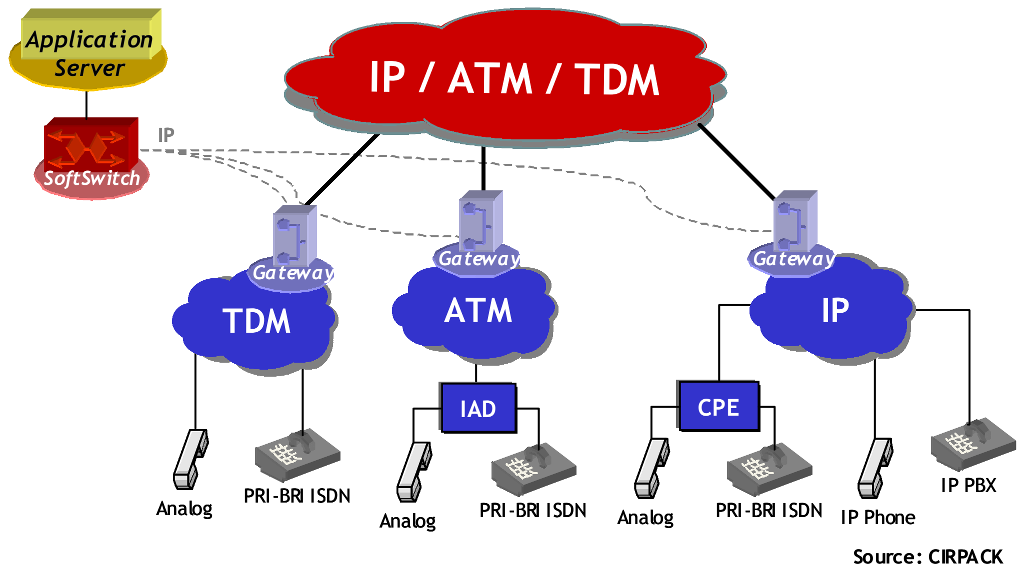

body responsible for the IP. SIGTRAN defines a suite of protocols (see *Figure 3*) enabling IP networks to transport PSTN signaling, such as ISDN (Q.931) or SS7 messages (IDSN user part [ISUP], signaling connection control part [SCCP], etc.), between IP nodes such as gateways, softswitches, IP–based databases, etc. It is a critical protocol to allowing the NGNs to support and interoperate with legacy services and equipment. The three main components of SIGTRAN are simple control transport protocol (SCTP), M2UA, and M3UA.

SCTP is a new transport protocol designed as an alternative to transmission control protocol (TCP) and user datagram protocol (UDP), aiming to carry signaling messages reliably over IP networks. M2UA and M3UA are two of SIGTRAN's adaptation sub-layers, supporting specific primitives required by a particular signaling application protocol. Other adaptation sub-layer protocols defined by the IETF include M2PA, SUA, and IUA.

H.248
In Geneva, on August 4, 2000, International Telecommunications Union–Telecommunication Standardization Sector (ITU–T) Study Group 16, which is responsible for multimedia services and systems, and the media gateway control (MEGACO) Working Group of the (IETF) agreed on a single standard for the control of IP/PSTN gateway devices: H.248.

The new standard supports gateway devices passing voice, video, fax, and data traffic between conventional telephone and packet-based data networks, such as commercial IP networks or the Internet.

The unified support of H.248, from the IETF and ITU–T, facilitates scalable and seamless implementation of services and applications over circuit-switched and IP–based networks. As a result, H.248 now dominates gateway control protocols.

ISDN Transport over Packet
The association of H.248 and SIGTRAN enables ISDN signaling to be transported to the softswitch over packet networks (see *Figure 4*).

SS7 Transport over Packet
The association of H.248 and SIGTRAN enables SS7 signaling to be transported to the softswitch over packet networks (see *Figure 5*).

FIGURE 3

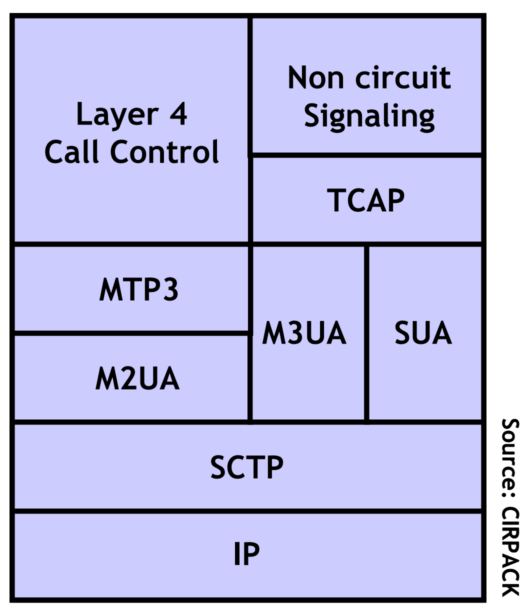

SIGTRAN Protocol Stack

FIGURE 4

SIGTRAN allows transparent transport of ISDN signaling over IP, and H.248 is the protocol controlling the gateway.

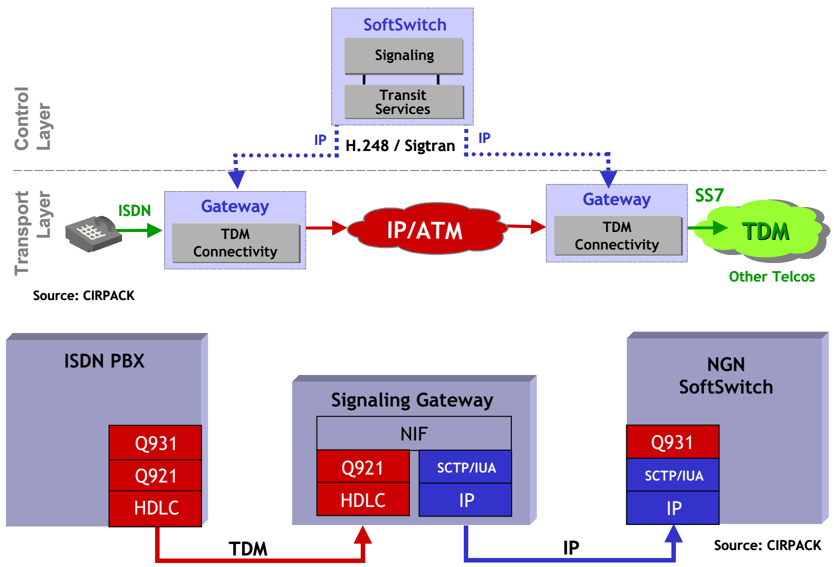

Signaling between Softswitches

The inter softswitch signaling must also support transparent transmission of SS7. SIP for telephone (SIP–T) enables ISUP messages transmission between softswitches.

SIP–T provides ISUP message transmission transparency across PSTN/IP connections. Traditional telecommunications services such as call waiting and 800 numbers are supported via SIP–T ISUP message translation.

Handling Pure End-to-End IP Telephony Signaling

Session Initiation Protocol

Session initiation protocol (SIP) was developed by the IETF and is a text-based control protocol for creating, modifying, and terminating multimedia sessions on an IP network, such as multiparty Internet conferencing and Internet telephony sessions. It supports real-time communications from simple audio calls to complex multimedia sessions. SIP supports unicast and multicast capabilities, is independent

FIGURE 5

SIGTRAN allows transparent transport of SS7 signaling over IP, and H.248 is the protocol controlling the gateway.

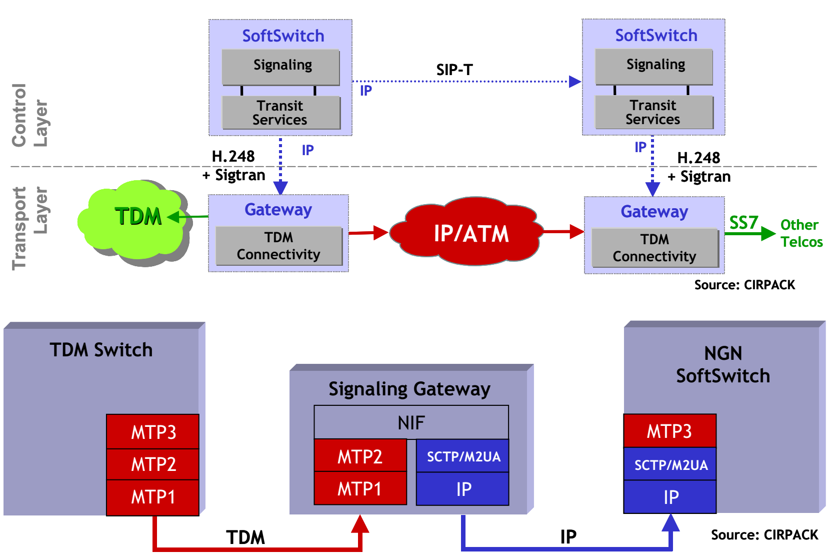

of the lower-layer transport protocol, and can be readily extended with additional capabilities.

In an IP telephony call, SIP can locate the called party and determine its capabilities. SIP supports mobility by creating proxies and redirecting requests to the location of the user. This allows the implementation of intelligent network (IN) features, such as name mapping, call forwarding, and call redirection. SIP address messages are formatted as human-readable text, which is compatible with the Internet address message format—also human-readable text.

With its scalability, flexibility, and performance characteristics, SIP is a leading contender for carrier-quality public IP–based telephone networks.

Media Gateway Control Protocol
Media gateway control protocol (MGCP) is a device control protocol that provides the means to interconnect a large number of IP telephony gateways, allowing them to work together as one (see *Figure 6*). MGCP is not an industry standard, and, due to a number of factors including its limited voice over IP (VoIP) networking functionality, is not expected to become one. However, MGCP eliminates the need for complex, processor-intensive IP telephony devices, simplifying them and lowering their cost.

MGCP and SIP are used by the softswitch to perform call control to IP telephony devices. Using MGCP, SIP, and H.248 to transmit analog phone traffic over a packet network, softswitch can also control the customer-premises equipment (CPE) or gateway.

Universal Telephony Signaling
By implementing a mixture of SIP, SIP–T, H.248, and SIGTRAN (plus MGCP in some cases), telecom operators can support all circuit and packet-telephony applications for perfect interoperability of their infrastructures with old and new services and networks of other providers (see *Table 1*).

4. Convergence Ahead

NGNs are gaining momentum in the telecommunications field. Not only do NGNs allow universal telephony, but

FIGURE 6

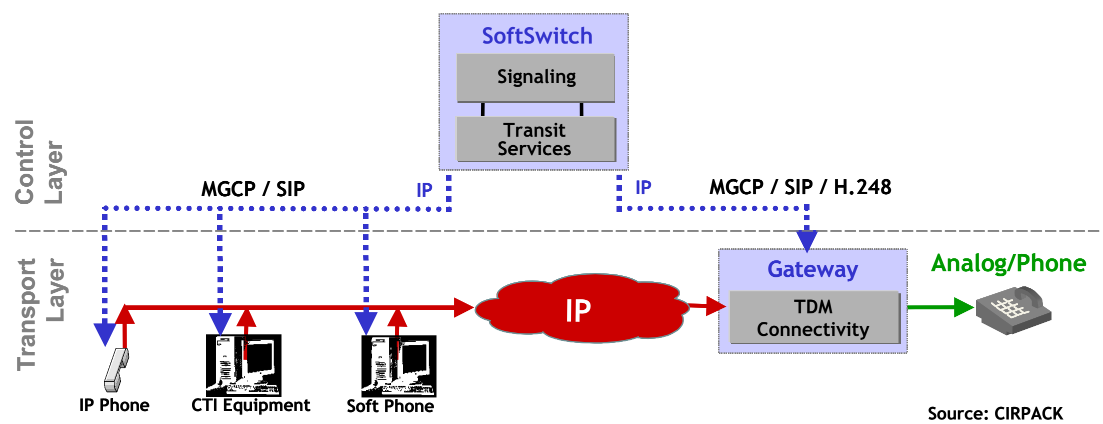

TABLE 1

Signaling for Use between Edge Device and Softswitch	SIP	MGCP	H.248	H.248 + SIGTRAN
IP phones	Best Fit	Best Fit	Feasible	
Soft phones	Best Fit		Feasible	
IP CPE with analog phones		Best Fit	Best Fit	
IP CPE with ISDN BRI phones				Best Fit
IP CPE with ISDN PRI link				Best Fit
IP Gateway with analog phones		Feasible	Best Fit	
IP Gateway with BRI phones				Best Fit
IP Gateway with PRI link				Best Fit

Source: CIRPACK

they are also essential for service providers in order to leverage their data networks and to increase profitability by offering innovative converging services. Powerful protocols are now available to ensure perfect interoperability between legacy circuit-switched public networks and NGN packet-based infrastructures and to maintain service continuity between users connected to various local loops and operators.

Operators that have massively deployed TDM technologies can move to NGNs while preserving investments and the customer base. New entrants, such as Internet service providers (ISPs) owning packet infrastructures, can now easily offer a full range of voice services at a fraction of the cost by deploying signaling gateways and Class-5 softswitches. The race to NGN is now open, as it allows the introduction of new breeds of services to catch the minds of subscribers—and their money. The current success of broadband Internet access, and the inevitable rise of local-loop unbundling, will accelerate deployment of NGNs.

The Future of IP Services

Greg Howard

Principal Analyst and Founder

The HTRC Group, LLC

1. Executive Summary

Internet protocol (IP) services represent an increasingly large opportunity for service providers, information technology (IT) companies, and enterprises. Yet the expense and inflexibility of current network solutions continue to constrain the market for IP services.

Network service providers (NSPs) continue to witness an explosion in data traffic. Yet generating high-margin revenue and creating unique customer value from selling pure access and bandwidth remains a challenge. NSPs are examining how to break out of traditional access service models and into high-margin IP services.

Application service providers (ASPs) have been offering IT outsourcing to enterprise customers for some years. However, this form of outsourcing often requires expensive private or leased network arrangements to meet network quality needs. Although the Internet provides a cheaper delivery alternative, performance and security are uncertain.

Enterprise customers desiring advanced IP services typically make expensive capital and networking investments, limiting the penetration of IP services across the industry. Competitive companies continue to focus on their core and strategic competencies while outsourcing IT applications and services. Several factors, however, constrain the outsourcing trend, including service flexibility, concerns over service quality and security, and expensive network access arrangements.

All consumers and suppliers of IP services will benefit from a service-delivery architecture that addresses fundamental business requirements. Network connectivity between ASPs and consumers needs to be as reliable and secure as private leased-line networks with the affordability of the Internet. In addition, enterprise customers want the flexibility to purchase and consume advanced IP services when required, without expensive IT and networking infrastructure.

A new architecture for advanced IP services will consist of partnerships between three suppliers: the service portal, the ASP, and the NSP. The architecture reflects the well-established business relationships between retailer, wholesaler, distributor, and end-customer.

Service portals are easy-to-navigate, point-and-click, Web-based interfaces that let customers order any number of advanced IP services—such as unified messaging (UM), software rental, or videoconferencing. The service portal is the retail storefront for advanced IP services.

ASPs sell wholesale advanced IP services to the service portal, which then retails the service to the business customer. ASPs focus on their core IT outsourcing competence and not on network services nor on marketing. With broadband service deployments solving bandwidth bottlenecks, ASPs will offer a broad range of attractive applications that will take the form of advanced IP services, such as voice over IP (VoIP) and videoconferencing over IP.

NSPs own the service networks that connect business customers to the service portals and ASPs. This distribution infrastructure allows consumers to obtain services that they have purchased from the service portal and that originate from the ASP. NSPs will deploy quality and security mechanisms in the IP network to offer secure applications across an affordable Internet infrastructure.

Through the examples of video and data conferencing, this document will show how all three players in the service-delivery architecture can profitably address the demand of enterprise customers for future IP services.

The result of the next revolution in IP services will be the delivery of outsourced applications to the enterprise market. NSPs, ASPs, and service [ortals will capitalize on the multi-billion dollar high-margin IP services revenue opportunity. The new architecture introduced in this paper provides a structure for delivering the next generation of IP services.

2. IP Services

Undeniably, business applications and services have gravitated to a common application and networking protocol—the Internet protocol. This section describes the current characteristics of business consumers and suppliers in this evolving IP services market, with the intent of articulating the key requirements of a future architecture to address the existing shortcomings.

2.1. Examples of IP Services

IP services are going beyond today's Web-site and e-commerce applications to provide advanced services of ever-increasing value to business around the world. Advanced IP services include the following:

- Integrated voice and data applications
- Videoconferencing
- Document storage and back-up
- Outsourcing of mission-critical applications
- Software rental
- UM
- Computer telephony integration (CTI)
- Network-based training services

Many of these applications and services are available in various forms today. However, their implementations are typically capital intensive and lack flexibility, resulting in limited deployment to a few early adopters. In addition, an inherent lack of end-to-end performance inhibits wide-scale deployments of these applications across the Internet.

2.2. Enterprise Customers

In the 1970s and 80s, a few large companies were able to achieve strategic advantage through enterprise-wide applications, but the implementation of IT strategies required costly investments in networking and data-processing technologies. While these infrastructures provided process automation across the enterprise, they were expensive to maintain and inflexible to changing business conditions.

Even as the costs of basic networking and applications continue to decrease, the costs of skilled IT staff to keep systems operating and properly maintained are increasing. Additionally, most chief information officers (CIOs) want to focus on the competitive advantage brought by information strategy and not the "bits and bytes" of information technology.

These trends have motivated many CIOs to focus capital and resources on core and strategic competencies and to outsource non-strategic IT applications and services. Companies no longer have to operate large and expensive data centers and private voice and data networks in-house.

This move toward outsourcing points to an even greater latent need: the unmet demand for guaranteed, secure, flexible, and affordable IP–based applications and services.

2.3. Service Providers

2.3.1. Application Service Providers

First-generation ASPs have been offering IT application outsourcing for some years. Many ASPs provide scale economy benefits to their customers by sharing the costs of expensive data-center mainframes, databases, and infrastructure across many customers. Hence, customers can save money through outsourced data-center arrangements, compared to operating dedicated data centers in-house. However, these initial offerings have been limited in their flexibility, as most arrangements required multi-year commitments. In addition, expensive dedicated networking arrangements have typically been required to meet service-level guarantees to each customer site. As a result, the expense and commitment associated with traditional outsourcing has essentially limited the market to mid-sized and large enterprises.

Web hosting ranks as one of the more common IP application services offered today. The explosion of Web-based content and services has permanently transformed the way businesses interact with their customers and partners. A notable advancement in Web hosting has been the re-architecting of large, mission-critical applications into so-called "thin-client" IP applications. Thin client applications minimize bandwidth usage through communications between client workstations and centralized application and data servers. Companies such as Arepa.com and Interliant AppsOnline have adopted thin-client technology and are already offering applications for rent over an IP network.

While ASPs are achieving success in offering services over expensive private networks or via existing best-effort Internet networks, the IP services industry will explode when quality guarantees can be offered across the Internet. The market will take off when a broad range of enterprise-class IP services is within the reach of small to mid-sized companies. Offering end-to-end guarantees for service quality is necessary to support advanced IP services. For ASPs to increase their revenues and market-share, future IP services will need to be "bundled" offerings of both network and application components. The offerings must be easy to purchase, with flexible billing and configuration arrangements. Business customers will subscribe to applications that are tightly coupled with reliable, secure delivery. Ordering and using the service will be as easy as placing a phone call.

2.3.2. Network Service Providers

NSPs continue to witness an explosion in IP data traffic. To address these demands, today's providers offer a plethora of data-network access and bandwidth options.

NSPs are starting to offer affordable virtual private network (VPN) services based on IP technologies. Most businesses understand a data VPN's fundamental benefits: facilitating communication with customers, suppliers, and partners, and reducing wide-area network (WAN) costs by using public networks rather than expensive leased-line networks. Businesses can use VPNs to connect employees, work sites, and external organizations (including outsourced data centers). VPNs based on IP technologies have significant cost benefits relative to dedicated ATM or frame-relay offerings. However, IP technologies are currently limited in their ability to support bandwidth guarantees.

While IP traffic continues to grow exponentially, generating high-margin revenue and creating unique customer value from selling pure bandwidth remains a challenge. Competitive pressures and technological advances have long since chiseled away the high margins on data access and bandwidth. Although NSPs are filling their networks with bandwidth due to ongoing improvements in broadband access technologies, the revenue produced from that bandwidth is relatively small compared to voice services and the margins are being eroded.

NSPs need to shift their focus to selling high-margin data-application services, analogous to the multitude of enhanced voice services (e.g., call waiting, messaging, voice-mail, voice conferencing, 1-800 numbers, etc.) that are offered to business customers along with basic voice transport.

To capture their share of future IP service revenues and margins, NSPs will need to overcome customers' concerns regarding IP bandwidth guarantees and security. NSPs will also need to offer on-demand configuration of IP VPN connectivity for network-based IP services to address the rapid, unpredictable pace of business today.

3. Requirements for Future IP Service Delivery

Business consumers and service suppliers will benefit from an infrastructure that supports the following common requirements: guaranteed quality of service (QoS), security, and on-demand subscription.

3.1. Guaranteed Quality of Service

In this paper, guaranteed QoS refers to the IP network's ability to deliver deterministic performance from end to end, measured by a number of specific parameters including bandwidth, delay, and jitter.

Standard IP QoS mechanisms (e.g., DiffServ) do not apply to specific applications or users, limiting the utility of these standard measures for granular control of mission-critical applications or for high-profile users. In contrast, guaranteed QoS will be delivered and measured by user and by application.

Another important distinction is that guaranteed QoS must be offered across the complete network connection—from the ASP to the business customer's local-area network (LAN). While ASPs typically offer service-level agreements (SLAs) within their areas of control (i.e., the data center) and NSPs typically offer SLAs between the edges of their core networks, what business consumers require is guaranteed QoS from the ASP data center to the business location.

3.2. Security

Security over IP networks has received a great deal of attention across the industry. Unlike frame relay, ATM, or private-line services, IP networks do not assign a "dedicated" physical or logical circuit between applications or users. Fortunately, the industry is making progress on security mechanisms across IP networks. Using tunneling technologies (IPSec), encryption algorithms (DES, triple–DES, Blowfish), key-exchange algorithms (Diffie-Hellman), and authentication mechanisms, security is being addressed. Business consumers can feel increasingly confident that sensitive business information can be securely transmitted across an IP network.

3.3. On-Demand Subscription

Another common requirement is the on-demand subscription, provisioning, and activation of future IP services. This capability allows the enterprise to respond quickly, with minimal cost, to new and unforeseen business needs. All the components of the architecture must facilitate this objective, including the dynamic provisioning of the appropriate guaranteed QoS and security starting from the ASP, across the NSP's network, to the enterprise consumer.

On-demand subscriptions must also be supported by usage-based billing mechanisms—much like pay-per-view movie services offered by cable television (CATV) operators. The cable network is similar to the IP network in that both are transparent to the subscriber. Customers will subscribe to advanced IP services and not to the transport and application infrastructure. Instead of purchasing bandwidth (e.g., 56K, T1, T3) users will be able to subscribe to a collection of services ranging from renting applications to scheduling videoconferences as needed.

4. A New Architecture for IP Services

The objective of this section is to provide a guide for service providers on implementing a new architecture to support the on-demand subscription of advanced IP services.

4.1. The Three Partners in the New Architecture

Three partners have unique roles in supplying advanced IP services to enterprise customers in the proposed architec-

FIGURE 1

New Architechture for Delivering IP Services

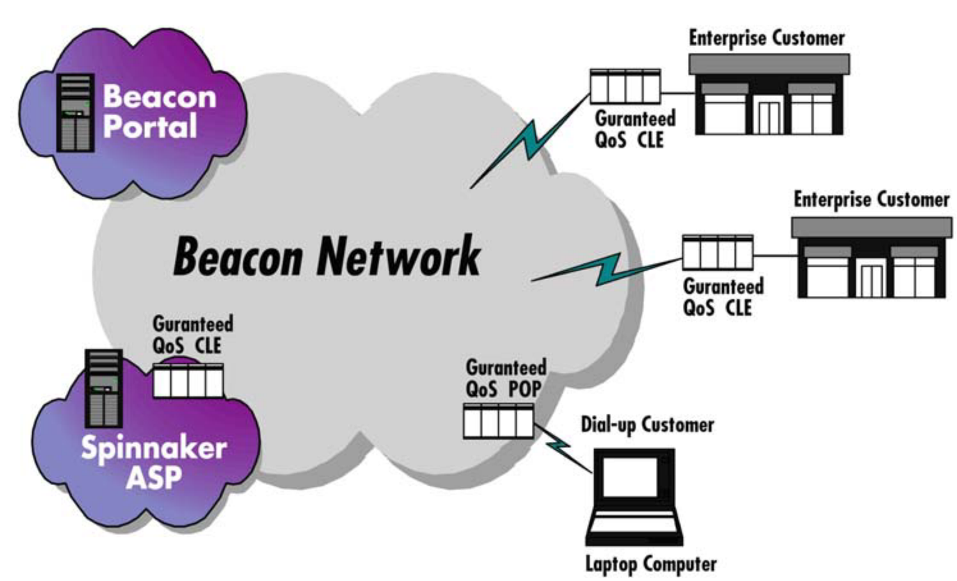

ture: service portals, ASPs, and NSPs. *Figure 1* shows the three partners working together. The architecture reflects the well-established business relationships between retailer, wholesaler, distributor, and business consumer. To illustrate the relationships, fictitious names have been assigned to the different elements. Beacon Network Services, Inc. is an NSP that operates both the service portal (Beacon Portal) and the service network (Beacon Network). Spinnaker, Inc. is an ASP.

4.1.1. Service Portals

As has been said, service portals are the retail storefront for advanced IP services. They are easy-to-navigate, point-and-click, Web-based interfaces that let business consumers order from a number of advanced IP services, such as UM, software rental, or videoconferencing. Service portals deliver advanced IP services based on the bundling of required resources from both application and NSPs. In addition, service portals provide the first line of technical support and customer service for advanced IP services.

4.1.2. Application Service Providers

ASPs provide advanced IP services to service-portal customers and act as wholesalers to the service portals. With broadband service deployments resolving bandwidth bottlenecks, ASPs focus their core IT competence on providing a growing range of applications, including VoIP and videoconferencing.

4.1.3. Network Service Providers

NSPs are the distribution component of the IP services architecture and own the QoS–enabled network infrastructure that connects business consumers to the service portals and ASPs.

To guarantee QoS for each application, a service demarcation point at the customer's location is recommended. The LAN/WAN boundary is a queuing point, making it the best place to perform policy management and to guarantee performance and security. Secure, guaranteed QoS will work best with a new class of customer-located equipment (CLE) at each business location. By managing and operating CLE for broadband services and point of presence (POP) equipment for dial-up services, the NSP can control security and guaranteed QoS features for a myriad of advanced IP services.

4.2. Service Portals

4.2.1. Business Strategy

Beacon Portal—the service portal—generates revenues by offering advanced IP services to enterprise customers. Beacon Portal's core competency is marketing IP services to customers. Beacon Portal's key challenge is attracting and retaining customers to new services.

Advanced IP services will consist of the bundling of two wholesale service components. ASPs, such as Spinnaker, will provide the application services while NSPs, such as Beacon Network, will offer the network-connectivity services. In essence, Beacon Portal takes advantage of an arbitrage opportunity between wholesale and retail pricing while adding value in exceptional service bundling, customer care, and billing/settlement capabilities.

To combat competition and minimize churn, Beacon Portal will differentiate itself through the creation of "sticky services"—advanced IP services with a strong potential for retaining existing business customers as well as attracting new ones. For example, a sticky service is an integrated voice and data application. A Web site can be enabled with IP telephony so that a customer browsing the site can talk "online" with a live service representative. The Web site would have a guaranteed QoS connection to service representatives in order for users browsing the Web site to communicate effectively. Beacon Portal may also focus on vertical markets, such as financial, medical, or manufacturing sectors.

4.2.2. Service-Portal Operations

Implementing a service portal requires the development of close business and technology relationships with business consumers, ASPs, and NSPs. Beacon Portal will need to integrate its service-order systems with its wholesale providers. Service portal operational systems will use open, standards-based technologies to best integrate with NSPs and ASPs. The objective is to support on-demand subscription to the application and network resources required for the service.

In this example, business consumers order conferencing services on-demand through Beacon Portal's interactive, Web-based service portal. The service portal automatically contacts the NSP and the ASP to ensure service-resource availability. Once resource and service availability have been confirmed, Beacon Portal provides its customer with a service confirmation message and a password to allow access to the service. When the service has been activated, the conference participants can log on to the ASP for the videoconferencing call over guaranteed QoS network connections provided by the NSP. *Figure 2* shows Spinnaker providing services to enterprise customers.

4.2.3. Billing and Settlement Mechanisms

Service portal operations will also require integrated billing and settlement mechanisms with its wholesale service providers. Beacon Portal collects usage information from the network and bills the customer, retaining its fee and sending settlement notices to Beacon Network for the network resource used and to Spinnaker for provisioning the videoconference service. Providing a unified bill is a competitive advantage, solidifying the customer relationship. In the service architecture, Beacon Portal provides the unified bill.

Figure 3 shows revenue flow from the enterprise customer to Beacon Portal, which then pays settlement fees to Spinnaker and Beacon Network.

4.2.4. Managing the Customer

Beacon Portal is responsible for managing first-level customer-support operations. When a problem occurs, the business consumer first contacts Beacon Portal, which attempts to diagnose the problem through the integrated network-monitoring and trouble-management systems it maintains with its partners. If Beacon Portal's first-level support team cannot resolve the problem, they escalate the issue with the appropriate partner to take action and restore service.

4.2.5. Financial Example

Beacon Portal provides the retail storefront for advanced IP services. Beacon Portal offers Spinnaker's on-demand video- and data-conferencing services over Beacon Network's guaranteed QoS network services.

FIGURE 2

Conferencing Subscription

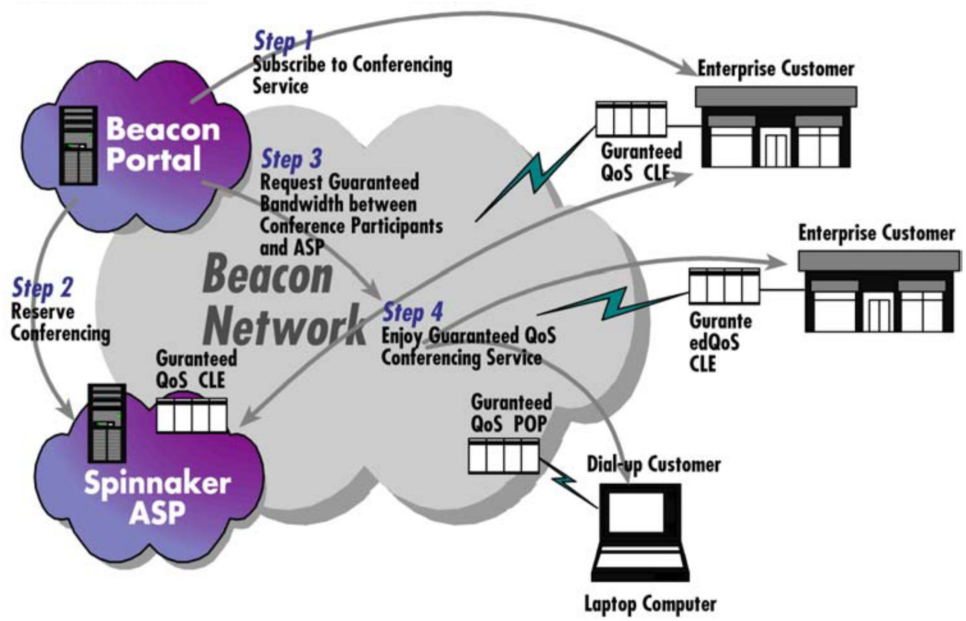

FIGURE 3

Billing and Settlements

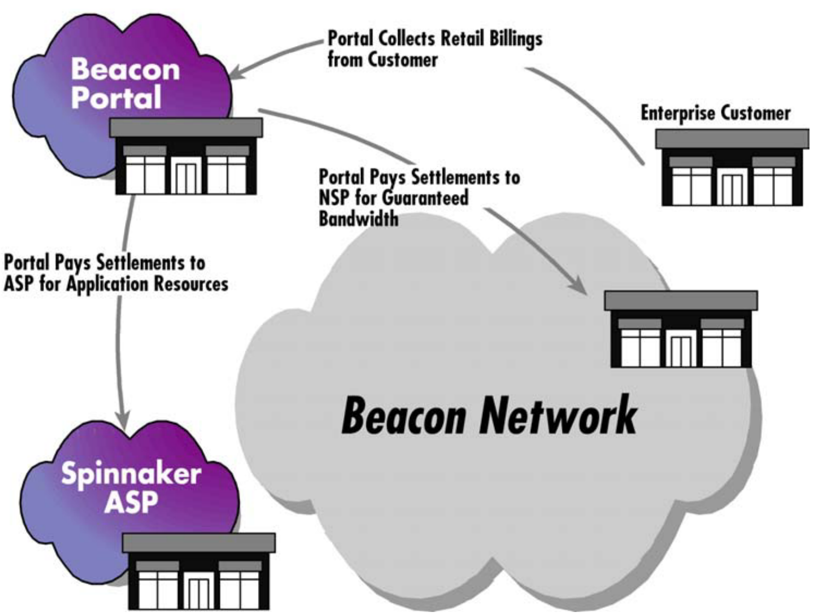

All of the financial examples presented in this document identify revenue, cost of sales, and gross profit over a three-year period. The assumptions used in the creation of all examples were based on feedback from service providers and portal companies, and on extrapolations from industry trends. The examples are intended to be used as a guideline only. Actual costs, revenues, and return on investment (ROI) may vary.

Common conferencing-service assumptions across all examples include an 8% adoption rate of existing customers being serviced by a large NSP. Customers use the conferencing service an average of 30 minutes per day.

Beacon Portal assumptions include the following:

- Retail revenue is based on charges of $0.35 metered rate per minute.
- Cost of sales includes settlement charges to Beacon Network and Spinnaker, support costs, and amortization of capital investments ($650,000 for Web Server,

customer care, on-demand service ordering, and billing and settlement systems).

4.2.6. The Bottom Line

As the retail storefront of the delivery architecture, Beacon Portal's revenues are substantial. Remembering that this financial example represents only one of what will be many services, the opportunity exists to generate significant revenues and cash flow.

The largest costs are the settlements paid to the ASP for the conferencing service and to the NSP for guaranteed QoS connectivity. After settlements, the largest operational cost is support, as Beacon Portal provides first-line support for its subscribers. The net margins, after sales, general, and administrative (SG&A) expenses following the first year, are around 8%.

Beacon Portal quickly recovers its initial capital investment of $650,000 for a billing and settlement system, customer-management software, and portal Web servers.

The cost of sales is high and the margins are relatively low compared to those of ASPs and NSPs, because Beacon Portal's business is an arbitrage opportunity between business customers and wholesale suppliers (ASPs and NSPs).

Like any retailer with wholesale suppliers, the service portal's business model is based on cash flow and volume. Beacon Portal has the largest revenue of the three partners because it is charging retail prices to end customers.

4.3. Application Service Providers

4.3.1. Business Strategy

ASPs specialize in providing network-based applications and services. In the example, Spinnaker partners with Beacon Portal to gain access to retail customers. Spinnaker generates wholesale revenue by delivering advanced IP services to the service portal's customers. Spinnaker relies on Beacon Portal to bundle and provision the guaranteed network connectivity for Spinnaker's advanced conferencing services. Spinnaker, a wholesaler of advanced IP services, can partner with many service portals to maximize its customer base.

Spinnaker's core competency is server-farm and application management. Spinnaker faces the challenge of building partnerships with application developers, portals, and NSPs. Application developers provide content and services for ASPs to host.

Historically, the relatively high cost and commitment associated with traditional outsourcing has limited the market to mid-sized and large enterprises. By lowering these barriers, the new architecture opens ASPs to a wide range of customers, from small businesses to large corporations.

PictureTel provides a real-life example of a developing ASP business in the context of the proposed architecture. PictureTel's IP–based videoconferencing application is being used by Enron Communications to deliver conferencing services to users. Enron is both the NSP and the service portal. PictureTel plans to offer its suite of conference applications through Enron to enterprise and service-provider customers.

4.3.2. Service Requirements

Applications may require that certain service characteristics, such as bandwidth and latency, be in place to fulfill the business customer's expectations. Spinnaker defines the service and specifies the required IP security and guaranteed QoS parameters required to support the service across the network. This is a one-time activity performed and negotiated with the service portal. Once the service definition is established, the service portal can support thousands of subscriptions to a single definition.

ASPs can offer a range of service-quality guarantees. For example, an ASP could specify three levels of service quality and security:

1. Real-time priority would be required for latency-sensitive applications such as voice over packet or videoconferencing. This quality level requires a low level of latency (e.g., less than 60 ms).
2. Mission-critical priority could be for applications less sensitive to latency but requiring strict security and encryption requirements. ASPs can assign a mission-critical quality level to enterprise resource planning (ERP) applications. Distinguishing ERP traffic from best-effort traffic guarantees that large data back-ups and file transfers do not interfere with ERP processing at critical times.
3. Best-effort priority could be assigned to applications such as data back-ups or e-mail. This level of service would be the most economical for non-latency dependent applications.

Spinnaker also specifies the parameters required to provision the service (e.g., IP addresses), how to bill for the service (e.g., $/minute of use), and what guaranteed bandwidth is required across the network (e.g., 356 Kbps per conference attendee).

TABLE 1

Beacon Portal Buisness Model

($000s)	Year 1	Year 2	Year 3
Revenues	$11,464	$63,086	$146,421
Cost of Sales	$8,987	$48,482	$112,240
Gross Profit	$2,477	$14,604	$34,181
Percentage of Revenues	22%	23%	23%

Once the service definition is agreed upon with the service portal, the integrated on-demand order processing software will reference this definition. This software system will forward the parameters for quality and service requirements, along with user-defined variables (e.g., IP addresses) to the ASP for application resources and to the NSP for IP network connectivity.

4.3.3. ASP Operations

ASPs host applications and make them available to business users. ASPs may develop their own applications and application content, or they may partner with application developers. Most of the ASPs provide IT hosting and outsourcing. Some ASPs also venture into software development, converting existing applications, such as ERP, to function as networked interactive applications. IT outsourcing companies see an opportunity in adapting to an on-demand provisioning business paradigm. And application developers view the ASP business as a new way to distribute applications to users, targeting new market segments facilitated by developing network technologies.

In the example, Spinnaker has established integrated provisioning, trouble-management, and settlement systems with Beacon Portal. Beacon Portal's relationships with the Beacon Network provides Spinnaker with the infrastructure that enables services to be delivered with true end-to-end QoS and security guarantees, extending service-quality commitments beyond Spinnaker's data centers.

4.3.4. Managing the Applications

Service problems that arise are initially identified and resolved by the service portal. When required, complex application problems are escalated and subsequently resolved by Spinnaker. To identify, track, and facilitate the resolution of problems, Beacon Portal integrates network-management systems and trouble-ticketing applications into the ASP infrastructure.

Customer-billing information is sent from Spinnaker to Beacon Portal, allowing the service portal to track usage and billing information.

4.3.5. Financial Example

Spinnaker, Inc. assumptions are as follows:

- Wholesale settlements are based on a $0.15 metered rate per customer minute from Beacon Portal.
- Cost of sales includes support and amortization of capital investments (conferencing servers, CRM software, billing and settlement systems, integration, and data-center facilities).

4.3.6. The Bottom Line

The wholesale revenues for Spinnaker's video- and data-conferencing services are significant. Remembering that this financial example represents only one of what could be many services, the opportunity exists to operate a profitable ASP.

Costs are dominated by the amortization of capital investment. In the example, Spinnaker makes a $3,750,000 investment for server-farm facilities, servers, conferencing software, billing and settlement systems, and integration. This initial investment accounts for the lower gross margin in the first year. Ongoing SG&A costs will likely be low because of the partnership with Beacon Portal, being that the service portal will bear most of the marketing and customer-acquisition costs, and will sell services directly to customers. Additional ASP costs will include developing new or enhanced application service offerings and building partnerships with application developers, portals, and NSPs. In the financial example, the net margins are 33% in the second year and 36% in the third year.

The financial model indicates that ASPs can expect significant margins on their revenues.

4.4. Network Service Providers

4.4.1. Business Strategy

NSPs provide on-demand IP network-connectivity services between business consumers and ASPs based on the security and guaranteed QoS parameters specified by the ASP for their particular service. Beacon Network provides access and bandwidth services. However, the commoditization of transport services means that the NSP must seek to support advanced IP services across its networks to benefit from the higher margins that these services can command.

The new architecture for delivering IP services enables Beacon Network to partner with third-party service portals and ASPs to deliver advanced IP services. NSPs can offer network-based services, such as voice over packet and VPNs, as well as network services packaged with application services offered by ASPs, such as software rental or data back-up services. Advanced IP services further drive bandwidth and application usage, increasing guaranteed QoS bandwidth consumption and thereby increasing Beacon Network's revenues.

Beacon Network's core competencies are operating a large network and maintaining guaranteed QoS connectivity to a

TABLE 2

Spinmaker, Inc. Business Model

($000s)	Year 1	Year 2	Year 3
Revenues	$4,913	$27,037	$62,752
Cost of Sales	$3,687	$13,818	$30,200
Gross Profit	$1,226	$13,219	$32,552
Percentage of Revenues	25%	49%	52%

significant number of customers. Beacon Network's key challenge is partnering with as many service portals and ASPs as possible to maximize its share of connectivity revenues for future IP services.

An actual example of an NSP operating in the new architecture is Enron Communications. Enron is an NSP that delivers high-bandwidth products and services to its customers. Enron has overbuilt its network to provide performance and has integrated intelligence into the network to deliver advanced IP services. Enron can define security and QoS on a per-application basis and it has partnered with companies such as PictureTel to offer advanced IP services designed to attract new customers and retain existing ones. Enron Communications' advanced IP service offerings include delivering live television content, videoconferencing, streaming video, distance learning, network-hosted applications, and new forms of entertainment. Enron Communications offers these advanced IP services to both broadband enterprise and service- provider customers.

4.4.2. NSP Operations

When business consumers order advanced, on-demand IP services through Beacon Portal, the service portal forwards the necessary service-provisioning information to Beacon Network. For pure network-based services (e.g., on-demand IP VPN connectivity), Beacon Network will provision guaranteed QoS connectivity between the requested locations. For ASP–based services (e.g., conferencing), Beacon Network will provision guaranteed QoS connectivity between the ASP's data center and the consumer location.

The NSP will require an on-demand provisioning system that can automatically configure the QoS requirements into the CLE based on subscription requests from Beacon Portal.

The integrated systems among Beacon Portal, Beacon Network, and Spinnaker provide the foundation to deliver guaranteed end-to-end QoS and security.

4.4.3. Technology Alternatives for an NSP

To offer advanced IP services, security and QoS must extend to the customer's location and must assign network resources by user or by application. Providers are just starting to implement QoS technologies, such as multiprotocol label switching (MPLS) and DiffServ, on their backbones today.

The network for delivering advanced IP services will be transparent to the business consumer and may include a variety of technologies and protocols such as asynchronous transfer mode (ATM) and IP routing utilizing copper, fiber, and wireless communications services. Both IP routed networks and ATM networks have advantages and disadvantages. Operating on Layer 3 of the open systems interconnection (OSI) model, IP routing is more common and relatively easy to maintain. The IP routing QoS mechanisms, however, are not yet deterministic. Operating at Layer 2 of the OSI model, ATM QoS mechanisms are more evolved and deterministic. ATM, however, is more complex to operate and maintain.

Most NSP networks will likely include a combination of both ATM and IP routing topologies, which together must be equipped to provide end-to-end service guarantees for advanced IP services.

4.4.4. Managing the Service Network

The service portal will identify and resolve service problems that arise. When required, Beacon Portal will escalate complex network problems to Beacon Network. To identify, track, and facilitate the resolution of problems, Beacon Network integrates the service portal's network-management systems and trouble-ticketing applications into the NSP infrastructure.

Customer billing information is sent from Beacon Network to the service portal so that Beacon Portal can track all usage and billing information. Hence, integrated settlement systems will be required.

Because the NSP owns and manages the infrastructure that supplies IP services to end users, it can provide service-level reporting to business consumers, service portals, and ASPs.

4.4.5. Financial Examples

Beacon Network generates revenue in two basic areas: basic access to the guaranteed QoS network and fees for the transport of usage-based advanced IP services. For simplicity, guaranteed QoS and advanced IP services opportunities are examined separately in this model.

1. The first financial example examines Beacon Network's investment and return for the basic access service to the guaranteed QoS network, including CLE.
2. The second financial example examines the business case for delivering an advanced IP service over the guaranteed QoS network.

TABLE 3

Beacon Network Guaranteed QoS Access Buisness Model

($000s)	Year 1	Year 2	Year 3
Revenues	$5,525	$29,660	$68,142
Cost of Sales	$6,016	$19,843	$42,006
Gross Profit	$491	$9,817	$26,136
Percentage of Revenues	–9%	33%	38%

4.4.5.1. Financial Example—Guaranteed QoS Access
Beacon Network assumptions are as follows

- Beacon Network has an existing QoS–enabled network backbone.
- Revenue from basic access services to the guaranteed QoS network assumes 2,000 customer connections added over a three-year period.
- Revenue sources include recurring access charges to the guaranteed QoS network and CLE installation fees.
- The cost of sales includes guaranteed QoS–enabled CLE, on-demand service activation operations support services (OSS), monthly local-loop circuit and leased equipment, installation, and OSS provisioning integration.

4.4.5.2. The Bottom Line—Guaranteed QoS Access
The financial example is typical of a broadband IP access service. Although revenues are significant and growing, margins will continue to be squeezed. Continuing to offer only basic access and bandwidth services does not present the most compelling business case.

The cost of sales includes capitalization of CLE and local-loop equipment, and expenses for CLE installation and support. Costs also include the capitalization of first-year investments, including $5 million for initial billing, on-demand service-activation OSSs, and systems integration. The first-year loss is largely due to the cost of integrating critical operations and provisioning systems. In the model, the net margins after the first year will increase from 7% in the second year to 18% in the third year.

The key value of the service network is the foundation that is created for high-margin revenues from the transport of advanced IP services.

4.4.5.3. Financial Example—Connectivity for Advanced IP Services
Beacon Network assumptions include the following:

- Wholesale settlements based on an $0.08 metered rate per videoconferencing customer minute and a $0.03 metered rate per secured data-conferencing customer minute from Beacon Portal.
- The cost of sales includes guaranteed QoS bandwidth between ASP conferencing servers and customers connected to the service network, as well as support and network operations.

4.4.5.4. The Bottom Line—Advanced IP Services
Although wholesale revenues from guaranteed QoS transport are expected to be less than the wholesale revenues generated by the ASP, the revenues and margins for Beacon Network are still significant. It is important to remember that this financial example represents only one of what will be many advanced IP services transported across the network. In addition, the revenues created by the delivery of advanced IP services are incremental—over and above existing access and bandwidth income generated by typical MNSPs today.

The incremental costs of additional bandwidth provisioned across an installed network infrastructure are quite low. Support costs are low because Beacon Portal is handling level-one support. Ongoing SG&A costs will likely be low

TABLE 4

Beacon Network Business Model

($000s)	Year 1	Year 2	Year 3
Revenues	$2,138	$11,766	$27,309
Cost of Sales	$1,163	$4,602	$10,152
Gross Profit	$975	$7,164	$17,157
Percentage of Revenues	46%	61%	63%

TABLE 5

	Core Competency	Revenue Magnitude	Gross Profit
Spinnaker ASP	Information Technology	Medium	Medium (52%)
Beacon Portal (Service Portal)	Marketing	Large	Low (23%)
Beacon Network (NSP)	Networking	Medium for Access Large for IP Services	Medium (38%) High (63%)

because of the partnership with Beacon Portal. In the model, the average net margin is about 36% over the first three years. This net margin is much higher than the margin for pure access.

5. The Result

Table 5 summarizes the key findings for the service portal, ASP, and NSP in the proposed architecture and illustrates the tremendous gain for all players in terms of core competency, revenue magnitude, and gross profit.

Enterprise customers will leverage the latest networked applications in an on-demand, pay-as-you-go format, with little incremental investment. Advanced IP services offer a viable and attractive outsourcing alternative to internal IT initiatives that are capital and resource intensive. Small and medium-sized enterprises, which previously could afford neither extensive internal IT departments nor high-end outsourcing, now gain access to a range of valuable IP services.

The three architecture partners will capture an increasing share of corporate IT spending by addressing the common requirements of guaranteed bandwidth, security, and on-demand service activation. The combination of services created by partnerships generates customer loyalty and reduces churn.

The architecture allows each player to focus on its core competency in turning the vision of advanced IP services into a reality.

IP Intelligence in the Access Platform

Jean Huppé

Director, Network Technology Development
Bell Canada

For the purposes of this paper, Internet protocol (IP) intelligence shall refer to IP services, particularly IP Layer-3 routing functions. Access vendors are incorporating more IP Layer-3 functions into their platforms. Typically, these new intelligence services are centralized in routers or boxes and implemented in the network core. However, IP Layer-3 routing functions are being pushed toward the network edges as an increase in demand for them emerges. This paper examines whether IP Layer-3 routing functions should be kept at the network edge or integrated into access platforms. Furthermore, it will address alternative approaches for delivering IP intelligence in the access platform.

Functionalities

Access vendors have worked hard to develop basic asymmetric digital subscriber line (ADSL) functionality and good digital subscriber line access multiplexer (DSLAM) capability. They now seek to add value to those products by incorporating more IP Layer-3 functions into their platforms, such as broadband–remote access servers (B–RAS), multicast functionalities, and voice softswitch gateways. The reality is that access vendors consistently introduce new ideas to service providers. *Figure 1* indicates the core, edge, and access layers of IP intelligence.

The core layer refers to the area of main concentration, and the edge refers to peripheral central offices (CO). The access layer serves as a link to customers. More mature functionalities, such as B–RASs, have been used for a couple of years and can be deployed in the DSLAM or at least very near to it. Multicast functionalities are also mature but are not traditionally integrated into the DSLAM further than the B–RAS location because point-to-point protocol over Ethernet (PPPoE) already connects to the B–RAS. Other Layer-3 service functionalities are being developed, such as multiprotocol label switching (MPLS) and IP virtual private networks (VPN) for business customers. Those functionalities are still being implemented in the network core but are being pushed to the edge. Eventually other functionalities will include a voice gateway softswitch and a voice over IP (VoIP) softswitch. More mature technologies are easier to interoperate between vendors and platforms, and standards for these technologies are also more mature. *Figure 2* illustrates the maturity level of various technologies.

However, many leading-edge technologies are less mature and proprietary in nature. It is often difficult to interconnect these less mature technologies to multiple vendors, as interoperability is often sacrificed for the sake of differentiation and speed to market.

Multiple Access Platforms

Incumbent local-exchange carriers (ILEC) have already deployed a series of different platforms in their networks and will continue diversification to deliver IP services to residential and business customers (see *Figure 3*).

Multiple generations or vendors of the same technology type may also be deployed. However, problems do arise when multiple access platforms are used to deliver IP services to the customer. For example, there is a need for many different types of platforms, such as DSLAM, very-high–data rate digital subscriber line (VDSL) in a multiple-dwelling unit (MDU), residential fiber-to-the-home (FTTH), business passive optical network (PON), and direct fiber. Multiple access platforms are likely to have several access vendors per platform for relationship diversity—it is possible that platforms such as ADSL will have two or three vendors in the next 5 to 10 years. In such situations, vendors may work well with one type of platform but not with others, which will obviously affect performance. Over the long term, the same vendor may also end up with different generations of products in the field. Bleeding-edge technologies have limited features and have a potential of a few years. Leading edge technologies may never have all the necessary features but have a potential of three to five years in the field, and mature technologies have a potential of 5 to 10 years in the field. For example, service providers might already have two generations of DSLAM in the field, such as standard-density and high-density DSLAM. Some advanced IP functionalities will be supported on the high-density generation but not on the standard-density generation. Therefore, multiple generations of products could exist even within the same platform and with the same vendor.

Distributed Intelligence

The advantages to distributed intelligence include rapid treatment of quality of service (QoS) and functionalities that are closer to customers. However, the disadvantages of distributed intelligence include the necessity for interoperability between multiple vendors. For example, if Layer-3 IP intelligence is distributed into each of those platforms, the same basic functionalities must be provided for different

FIGURE 1

Introduction

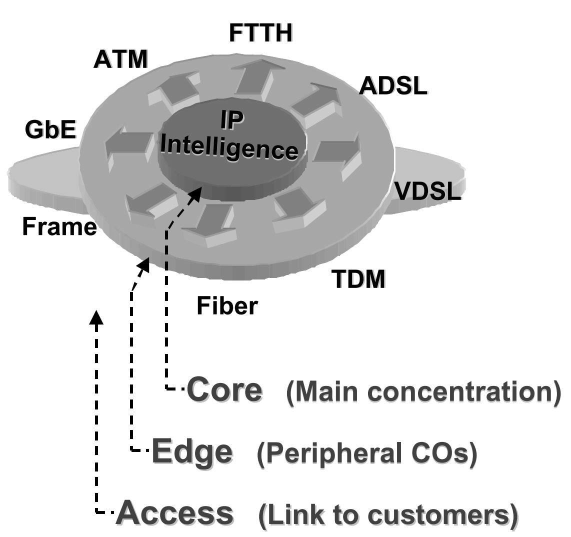

FIGURE 2

Prevalent IP Functions

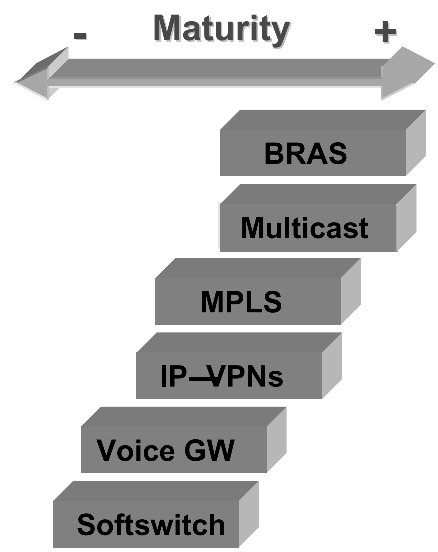

platforms, vendors, and generations of products. In addition, if an IP VPN service over MPLS is to be provided, it is essential to the customer that it looks and acts the same from any platform. *Figure 4* illustrates this point.

The disadvantage of distributed intelligence is that the availability of functions will not be the same. Each of those platforms indicated in *Figure 4* must be able to connect to different backbones in the network. For example, service providers with multiple backbones might have asynchronous transfer mode (ATM) for DSL, a multiprotocol label-switching (MPLS) backbone, and gigabit Ethernet (GbE) resilient packet rings (RPR). There are different access and different backbone platforms, which makes it difficult to fit those functionalities inside each of those platforms, an obvious disadvantage. In the end, distributed intelligence leads

FIGURE 3

Multiple Access Platforms

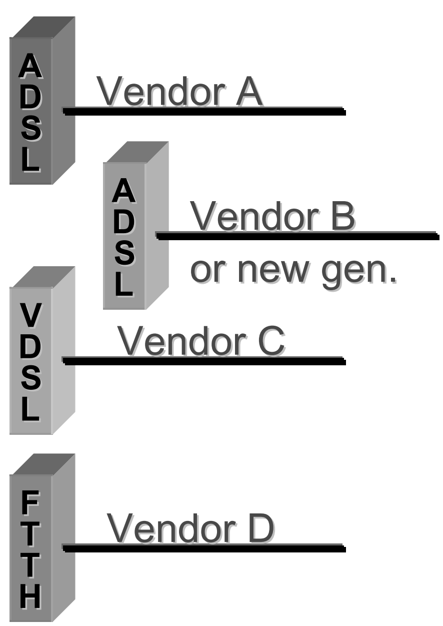

FIGURE 4

Distributed Intelligence

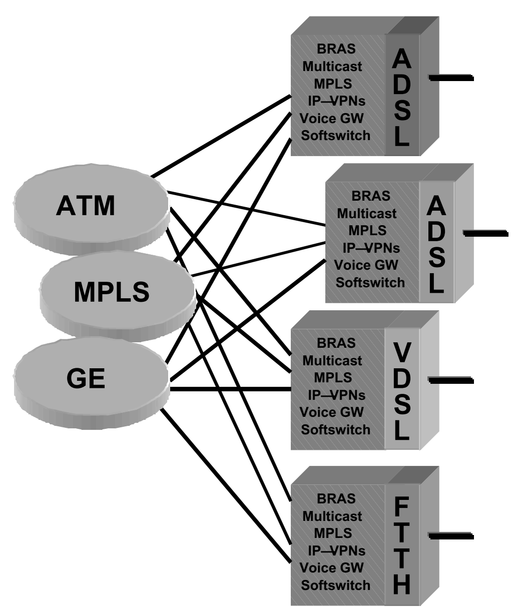

to more complex operations and more complex information services (IS) and information technology (IT).

Centralized Edge Intelligence
Another option would be to centralize those functionalities, placing them at the edge of the network and thus keeping them separate from the access platforms. This keeps those access technologies fairly independent of the Layer-3 functionality. The advantages of centralized edge intelligence include simplified management and uniformity of IP functions. It is easier to select a best of breed vendor of a particular IP function without being biased by the connectivity layer. Edge intelligence is independent of new access platforms, as it is not necessary to have functionalities connected to access platforms. Connectivity to the backbone network is thereby simplified (see *Figure 5*).

The main disadvantages to centralized edge intelligence include the addition of another box in the network and the potential for bandwidth bottleneck.

Super-Box Concept
Another solution would be to have a single box that could deliver everything from ADSL and VDSL to fiber to the home. It could also connect to all types of backbone infrastructures. This concept can be applied more easily in a competitive local-exchange carrier (CLEC) and data local-exchange carrier (DLEC) environment where the market is focused and people start from scratch. It is more difficult in an incumbent local-exchange carrier (ILEC) environment where there are a number of legacy access platforms delivering IP connectivity. *Figure 6* illustrates the super-box concept.

This super box does not function well in networks with multiple platforms, product generations, or vendors. Ultimately, even CLECs might have problems maintaining this super box in the network.

Conclusion

Good Layer-2 technology might be what is truly needed in the access platform. With limited bandwidth, it is necessary to have good QoS and class-of-service management capabilities. Mature functionalities, such as B–RAS or multicast, could be included in the access platform and not just beside it if they provide efficiency. In the access platform, it is advisable to avoid bleeding- and leading-edge Layer-3 functions that would incorporate proprietary developments by specific vendors unless it is the only access platform. An approach that is too decentralized will increase costs and should be avoided in the access platform.

FIGURE 5

Centralized Edge Intelligence

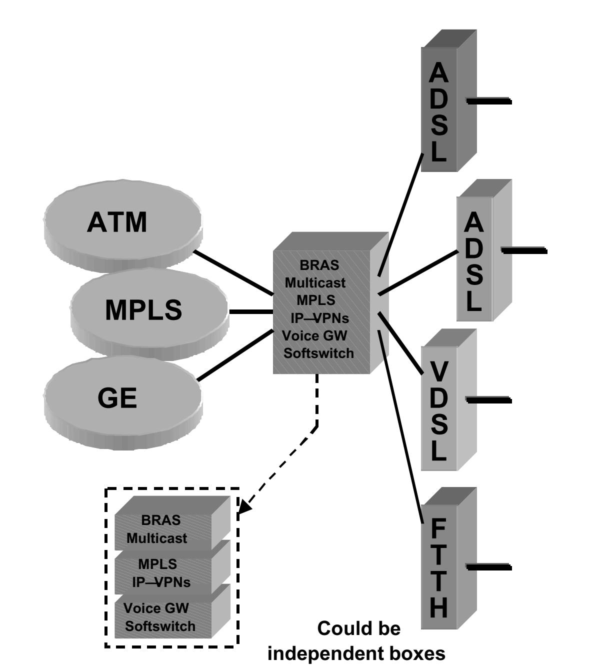

FIGURE 6

Super-Box Concept

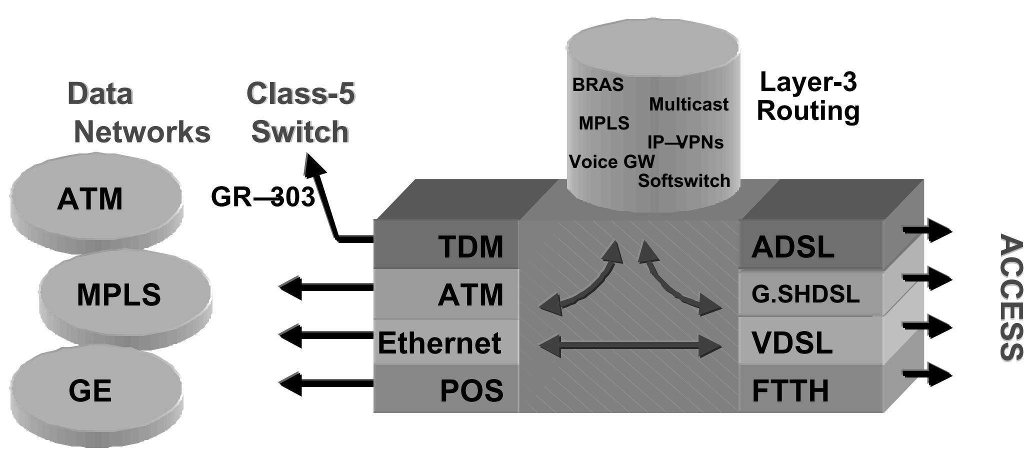

IP Telephony Interoperability: What It Takes to Be Carrier-Ready

Michael K. Kaczocha

Senior Engineer, Voice Evaluation and Research Team

AG Communication Systems

Introduction

Internet protocol (IP) telephony systems commonly claim interoperability with one another but don't go into much detail regarding what interoperability means. One of the main topics of conversation within the vendor community during the last few months has been the lack of progress with carriers and other customers in moving beyond non-revenue trials. These two are related. To have an offer that is ready for full deployment, it must perform as well as traditional equipment.

To verify that customer-premises equipment (CPE) devices provide the feature set, stability, administrative capability, and conversation capability required to become a recommended product for customer use on the iMerge Centrex feature gateway (CFG), the device must undergo a series of tests and real-life applications prior to being qualified. The tests and live-use applications performed on these products are broken into three phases: 1) iMerge CFG to CPE interoperability, 2) voice quality, and 3) internal employee trial.

The first two phases of testing are performed in a controlled laboratory setting where all aspects of the testing are under the control of the test person. The third phase of testing, the internal employee trial, involves subjecting the CPE device to real customers (company employees) by providing them the device for use in their normal day-to-day telephone activities. Using these three testing and equipment application environments when testing each of the qualified devices assures those customers who purchase the iMerge CFG and utilize the qualified CPE devices that the iMerge CPE solution will provide a level of service that satisfies all needs and expectations.

The CPE devices that undergo qualification testing are divided into five categories: 1) analog IP phones, 2) integrated services digital network (ISDN) IP phones, 3) analog access gateways, 4) ISDN access gateways, and 5) softphones. Each category of devices has its own set of test suites that are executed for that device type.

Prior to CPE vendors' submitting their device for qualification testing, documentation is provided that defines the expected operation of the CPE device when connected to an iMerge CFG. This document provides specific signaling/message-handling guidelines for the multitude of voice path set-up, configuration, and maintenance messages that can be presented to the CPE device over the packet network. In addition to the signaling characteristics document, the CPE manufacturers are provided a functional specification document that defines the expected operation of the device. These documents are used as the foundation for generating the test suites used during qualification testing.

Interoperability Testing

The first phase of testing undertaken is interoperability testing. Interoperability testing is designed to prove that the device can be configured and administered by a system administrator, that the device is capable of performing a predefined feature set, that it is stable, that it provides reliable operation, and that it can be maintained in a live customer environment. To prove these factors, the tests include the following:

- Registration capabilities
- Signaling capabilities (in-band, out-of-band)
- Q. 931 compliance
- Subjective conversation capability
  Place real calls from the device being tested to various network terminations (e.g., device being tested to a plain old telephone service [POTS] analog telephone)
- Verification that product aspects that influence voice quality and general operation have been addressed. In cases where traditional telephony standards exist, they are used.
  - Echo cancellation
  - Dial-tone delay
  - Ringing delay
  - Ring-back tone delay
  - Flash/on-hook delay
  - Differentiates services (DiffServ) tagging
  - Comfort noise
  - Packet loss
  - Packetization rates
- Verification of conversation capability at supported compression rates

- Basic telephony capabilities (e.g., flash, voice-mail message-waiting indicator, hold, three-way calling, distinctive ringing, speed dial, redial, etc.)
- Administrative capabilities
- Provisioning capabilities
- Security
- Test load simulation (for gateway type devices)

Using an automated call generator, it is determined how well the devices perform under repeated call originations and terminations. The performance of the device is measured according to its lost-call rate (For every 10,000 calls provided to the device, how many calls were completed?). IP endpoints are tested with an eye toward the same standards of performance as traditional endpoints.

Voice-Quality Testing

The second phase of testing involves evaluating the voice quality of a device. How humans perceive sound and, in turn, judge voice quality is not straightforward enough to allow voice quality to be easily or accurately measured solely with objective test methods. The voice-quality evaluation approach includes both subjective and objective testing, but overall emphasis is placed on subjective testing, since voice quality is subjectively perceived.

For subjective voice-quality testing, experienced/expert listeners, who are trained to listen for and identify impairments in voice quality, are used. This method has advantages in that it provides judgment as to general voice quality as well as providing additional information as to the actual impairments heard and surrounding circumstances. This also allows further investigation and problem resolution and provides the information needed to test updates to the product to validate impairment resolution.

Objective methods are used to measure an attribute or combination of attributes related to voice quality by using test equipment. Objective tests are good for pointing out potential problems and are excellent tools for analyzing impairments and their root causes.

The iMerge testing approach incorporates several individual objective measures, such as delay and loss, but does not rely solely on the objective predictors because of their known shortcomings and lack of informative results.

To ensure a user's perception of voice quality is addressed, voice-quality evaluation is separated into three testable components. These components are voice clarity, echo, and delay.

Because subjective conversational testing is the most straightforward and expedient way to obtain initial voice-quality measurements, in most cases the evaluation starts with this type testing. Also, there are some impairments that may only show up in conversation testing.

Examples of some of the voice-clarity problems identified in this phase of testing are as follows:

- Distorted speech
- Speech too loud
- Speech not loud enough
- Clipping of speech or noise (in double-talk situations)
- Clipping of speech or noise (without double-talk)
- Echo
- Awkward conversation interaction—delay
- Lack of noise, feels like connection dropped
- Extra annoying noise(s)
- Fluctuating speech or noise levels

An additional conversational test method used for evaluating voice quality and for reproducing impairments is the recorded conversational method (RCM). For RCM testing, digitally recorded two-way conversations between various types of voices (e.g., male-to-male, male-to-female, female-to-female) are played through the system, and the resultant listening path is digitally recorded. The digital recording of the listening path is then played back and analyzed for perceptible quality degradations.

The initial testing for determining whether echo is a problem is accomplished through the conversation tests mentioned in the voice-clarity section. Another method used is the composite source signal (CSS) echo loss test. This test is performed by playing a digital recording into the device under test, recording the return listening path and analyzing the recording for instances of echo. A G.168 echo loss test is also used for echo detection. This echo loss test varies slightly from the CSS test in that it tries to match the echo loss test mentioned in G.168 documentation.

To measure the delay component of voice clarity, periodic measurements are taken, and then an average is calculated. Test equipment that utilizes a recorded signal rather than a pulse to measure delay is the primary method used. For this, test a signal that contains many different properties of speech is delivered in a single recording. The recorded .wav file analysis method is also used for measuring delay. This test is performed by simultaneously playing a digital recording (pulse followed by speech) into the path, recording the listening path on both ends of the connection and analyzing the recordings using a software program.

Other factors that can influence voice quality and which are considered during voice-quality evaluation are as follows:

- Packetization rates
- Jitter buffer
- Codecs
- Voice packet loss
- Voice packet jitter

Internal Employee Use

The third phase of testing CPE devices involves subjecting the device to a real-life environment where real users (company employees) utilize the CPE device for everyday phone use.

In the case of analog IP phones and ISDN IP phones, the users normal desk phone is replaced by one of these devices, which is then used as their normal telephone. The IP phone is connected to an iMerge CFG and then to the Centrex switch, which provides their telephone services. For analog and ISDN access gateways, the normal desk phone will be routed to the gateway and then to the iMerge CFG, which is connected to the building Centrex switching system. For both the IP phone and the access gateway application, the user will continue to use the same telephone features that he or she has used up to the point of converting to the new connection.

Internal field trials for softphones are handled in a very similar manner. Employees are provided the software and hardware, where applicable, and they utilize the device in a direct local-area network (LAN) hook-up or in a remote dial-up environment. For remote dial-up applications, the user/employee is provided a connection to a building modem pool, which provides access to the iMerge CFG and the company Centrex switching system, thereby providing he or she with the same telephone services that he or she would have while working in the office. The Centrex switch used in the employee trial is a live switch that has all of the characteristics of a switching system that would be utilized in a live customer environment.

During the trial period, the end users are solicited to provide information regarding performance and satisfaction level. In addition to obtaining valuable feedback from real users, very important information regarding application guidelines (e.g., how to configure the device in a non-lab network) is obtained. The internal trial is a very important part of the testing. It provides the confidence that each CPE device and the features associated with that device can operate in a real-world environment.

Conclusion

By subjecting a CPE device to the comprehensive battery of testing discussed, a robust solution can be demonstrated to a carrier whose customers have high expectations of the reliability and quality of their communication systems. Many IP telephony systems have not met those expectations; they have been plagued by poor reliability, limited functionality, and disappointing sound quality. Carriers are looking beyond the hype attached to voice over IP (VoIP) and are measuring the performance of these products to the same standards as traditional systems. A testing regimen such as that described provides assurance that the “table stakes” attributes of reliability and quality of service (QoS) are available.

Many vendors will further evaluate equipment on human factors and other traits, which are also critical in the carrier arena. These do not fall under the realm of interoperability and are not addressed in this paper, but a vendor would do well to look beyond the scope of this paper if they expect to achieve broad deployment by a carrier.

The Dilemma of Digital Transition

H. Steven Klein

Vice President, Business Development

Tut Systems

1. Introduction

1.1. Problem Statement

As the Post Telephone and Telegraph Administration (PTT), regional Bell operating companies (RBOCs), and telecommunications operators begin offering digital video services, they are faced with some key economic and technical obstacles in trying to create a competitive model against embedded cable television (CATV) service providers. With two-way digital cable systems, the fixed capital cost to pass each home is very high in relation to customer-premises equipment (CPE) and in-home installation costs. On average, infrastructure capital costs are about two-thirds of the cable operator's installed cost percentage, and the combined CPE and in-home installation costs are about one-third. This is due to the cable operator's need to pass every home, whether or not service is taken.

The phone-company service provider looking to enter the video arena and compete with cable operators using digital subscriber line (DSL) technology face the opposite problem than that of the cable operator. The nature of DSL access systems allows system capacity to be engineered to a planned take rate rather than pass every home. Therefore, the telco service provider has low fixed capital costs to build the DSL network to support video.

However, all video and data services in the home are digital, unlike a two-way cable system that has both analog and digital channels. Therefore, a set-top box (STB) is required at each television. According to recent surveys, the average U.S. home has 2.3 television sets. This implies that in order to offer a complete video service, telco operators need to provide three STBs to a majority of homes. And the various distribution schemes in the home to distribute services digitally to appliances often require extensive in-home wiring. The end result is the telco operators sees one-third of its capital costs in fixed infrastructure and two-thirds in the variable CPE and in-home wiring costs on an installed-cost basis.

Because the cable operator has such high fixed capital costs for infrastructure, take rates and churn become critical factors to profitability. Cable operators focus a great deal of time and effort on capturing and keeping as many customers as possible. In the past, the typical 15 percent annual churn experienced by cable operators has had a direct and negative impact on profitability.

On the other side of the coin, telco operators must enter a market already served by cable operators and offer competing services with hopes of taking customers away from the incumbent. The low infrastructure capital costs favor them in that regard. But once the customer is won, they must absorb a higher cost to place STBs in the home and perform installation.

This paper looks at the issues and problems related to CPE in the home and at the in-home wiring options for the telco operator offering a bundled video and data service using DSL. It examines the issues, as well as the alternatives and solutions, as they relate to making a successful business case. It attempts to identify the most economic ways to deliver a whole home solution with DSL.

1.2. CATV Background

Until recently, CATV systems were analog only, and therefore did not require an STB at every television in the home. An STB was used only if the subscriber purchased premium channel packages that required descrambling, and cable operators could charge a fee for the services as well as for the rental of the STB.

Today, cable operators are upgrading their systems to two-way digital and now require a STB to access digital broadcast. But, because the hybrid fiber/coax (HFC) cable system remains both analog and digital, the subscriber can have digital on selected television sets (with an STB) and remain on analog for secondary televisions with no STB. Likewise, subscribers can choose to purchase more expensive digital broadcast service packages, or not. Therefore, only those customers upgrading from analog to digital need an STB. In the United States, about 15 percent of the entire cable subscription base uses an STB today.

1.3. The Dilemma for Telcos

Telcos offering video over DSL must provide a set-top box to all subscribers and all televisions. They do not have the option of offering analog to a specific subscriber or certain televisions within the living unit as cable operators do. So the initial capital investment in CPE devices supporting digital video services over DSL is considerably higher than that faced by cable operators.

Over time, cable operators will increase their digital footprint and service penetration and will also have to provide an increasing number of CPE devices. In the United States, numerous cable industry surveys assert that the average number of televisions per household is in the 2.3 to 2.4 range. Already, AT&T Broadband offers two digital STBs in its Platinum Service Digital Package, in recognition that more than one television exists per home. Forecasts for digital STBs in the CATV industry indicate that two-way digital penetration will reach 54 percent (50 million households) across all U.S. households by the end of 2005. This will create a more even parity in competitive service offerings with telcos, as well as closer parity in CPE costs per subscriber.

1.4. Implications to Telco Video Providers

Early DSL business models (harkening back to Video Dialtone in the early 1990s) assumed a value-added service overlay to existing CATV services. In this model, DSL would be used to deliver digital video to a STB located on one television (the main television), with a service offering including video on demand (VOD), specialty channels, and interactivity/games. Rather than competing with cable, the telco would offer its own unique and differentiated services. Over time, this model has proven ineffectual. Current thinking is geared to offering "full service" video directly competing with cable plus enhanced services as part of the bundled full service offering.

Telcos that must compete with the cable operator, which implies taking the customer away from the incumbent cable provider, must also provide service to more than one television in the home where this situation exists. There is no business case rationale for a subscriber to keep his analog cable service and also purchase digital service for a selected television from another provider. Competitively, it becomes an "all or nothing" proposition.

It should be noted that the models differ around the world based upon governmental policy, etc. Therefore, a series of business models exist imposing design and development challenges to equipment manufacturers trying to come up with solutions for video over DSL. The notion of "one box fits all" simply does not play out.

Several different approaches are under consideration to address serving televisions in the home. The biggest driver, outside of the applications and service requirements, is economics. A combination of lowest feasible electronics (CPE) costs, and practical matters such as ease of installation, all come into play.

2. Pros and Cons of Various CPE Models

2.1. Analog and Digital RF Combining at the STB

In a number of European countries, the government requires service providers to offer "free to air" channels as basic broadcast. There is little interest by telcos to absorb the cost to encode and support free to air channels for which they get no revenues. At the same time, a telco cannot offer its pay-TV services because doing so would require it to disconnect the subscriber from access to free to air content.

A solution to this issue would be radio frequency (RF) combining of analog and digital over the in-home coax, with the STB providing both a digital decoder and analog RF output. This type of solution would resolve the issue posed in the analog + digital CATV model, where a need exists to provide both types of content (channels). And it allows the telco operator to offer its premium pay-TV digital services to a main television in the living unit. There is an added cost for the RF in combining technology at the STB, but it is probably much more acceptable than other alternatives. However, this solution is probably not practical in the United States, where even analog cable service is a pay service. The telco operator would not want to provide both analog and digital services nor would the subscriber purchase two services from different providers.

2.2. Distributed Set-Top Box

The FS-VDSL Committee is a group of more than 70 international companies that share a common vision for an end-to-end multiservice, video-centric network platform based on very-high–data rate DSL (VDSL) technology. Many members of the group, and others around the world, favor an open DSL interface at the home, with Ethernet or HPNA 2.0 (a de facto industry standard endorsed by the Home Phoneline Networking Alliance, or HomePNA) used for in-home distribution to discrete STBs at each television to be served. The principle behind this is to segregate access from CPE and eventually open the path for set-top devices to be purchased by consumers through retail distribution. It may also be a "mandatory" requirement to have a service separation point for regulatory reasons as well as to allow an operator to support both wholesale and retail services from multiple service providers.

The opposing viewpoint is a single residential gateway device with multistream Moving Pictures Experts Group (MPEG)-2 outputs, requiring a single piece of CPE in the home for video and data service connectivity. The distributed model requires a modem device with routing capability to distribute bandwidth to each discrete device in the home, as well as an STB on each television served. In today's environment, the economics are significantly higher in the distributed model.

A major argument, in favor of the distributed model, is the hardware costs associated with offering high-end interactive television services versus basic broadcast service. First, the digital service suite continues to expand, including VOD, games, and other interactive applications, as well as personal video recorder (PVR) functionality. There is a cost associated with such applications in the form of hardware and client software. In a distributed model, a feature-rich STB could be placed at the television, where these services are required (and additional fees are paid for), while other televisions could be served with low-cost basic STBs. In a residential gateway model, the cost for all of the additional hardware and software is paid for in the box for all televisions, even if the subscriber is only willing to pay for the enhanced features for one television.

Second, the distributed model allows operators the ability to charge for services on a per-TV basis. In the current residential gateway business model, the customer pays the same for digital service as cable charges but gets the capabilities enabled to every television (something cable does not do without extra fees).

Third, and importantly, the distributed model allows CPE to migrate to retail distribution. This gets the operator out from under the capital cost burden of providing the CPE to enable service. This type of model is already established in the direct broadcast satellite (DBS) and TiVo/Ultimate TV markets.

2.3. *Residential Gateway Model*

The concept of a single CPE ("mega STB") residing at the main television and providing multiple MPEG-2 video stream and data connectivity has enjoyed commercial deployment for a couple of years. This solution leverages common electronic components on a single board to provide what is in effect the functionality of three STBs and a modem in a single unit. Compared to CATV digital STBs (such as the Motorola DCT2000) or DSL single STBs, the cost of the residential gateway is significantly less. Assuming that the average of 2.3 televisions per household is accurate, a three-stream residential gateway can solve the "digital dilemma" in a large percentage of U.S. homes.

While the economics appear to favor a residential gateway model over the distributed CPE model, there are some significant issues and downsides to be considered, such as the following:

- Currently, the only commercially offered residential gateways are proprietary to a specific vendor's access system, not allowing their use with other access systems or technologies.
- In households with only one television (much more common in Europe and the Far East than in the United States), a residential gateway is actually more costly. Having multiple televisions to serve in every home is needed to justify the residential-gateway business model.
- There are additional "peripheral requirements" to enable a residential gateway to work, such as antennas for receiving remote control signals from television locations away from the main television where the residential gateway resides. *(Wiring issues are discussed later.)*
- Unless the residential gateway becomes a "generic" interface and non-proprietary, it can never migrate to a retail distribution model.
- Whenever hardware and software enhancements are added (such as a PVR hard drive, high bit-rate graphics for gaming, etc.), the price of the residential gateway increases, even if the features are only needed at one television. The box cannot be feature-scaled to the service needed at each television based on its integrated design.

2.4. *"Combo Box" Model*

Bell Canada's ExpressVu is promoting a combo box model, where digital video is input via dual DBS receivers, and asymmetric DSL (ADSL) technology is used to deliver VOD, data, and interactive services. It assumes DBS is more cost-effective than VDSL in a residential single-family home serving area, and the longer rate and reach that ADSL affords allows a data and interactive overlay to broadcast from the satellite.

The counter arguments to a combo box approach include the costs for the satellite dish and installation, the cost of the box itself, and the limited use of DSB in high-rise multidwelling unit (MDU) buildings. It will have to be seen if this approach is commercially successful (or the product actually works as specified).

2.5. *Conclusions*

The biggest issue today is not that of product quality but of volume. DSL–based set-top and residential-gateway boxes are high volume in a relative sense because there has yet to be any mass-scale deployment of video services over copper by telcos. Whether it is Pace, Thompson, Motorola, or NextLevel Communications, it is hard to leverage reduced component costs and manufacturing efficiencies when annual volumes are sub-100,000 worldwide for set-top devices. Cable, on the other hand, can leverage manufacturing volumes into the millions for STBs. It will take a worldwide demand of one million or more to begin seeing volume reflected in pricing. Also, it will take that kind of volume to engage large consumer electronics manufacturers and retail outlets to enter the game.

Secondly, through the efforts of worldwide consortiums, such as FS-VDSL, the industry must be driven to some common ground regarding CPE reference designs and open interfaces. One of several things will need to happen to allow leveraging of volume across a large telco market segment worldwide:

- Either the proprietary interface to the residential gateway must change to an open one to allow common product support across various access systems (assuming the worldwide consortium agrees that such a residential gateway concept is the best solution overall); or
- The consortium defines a common interface and signaling methodology that allows a residential-gateway design to be licensed and manufactured by consumer electronics companies for the worldwide market (including retail); or
- The worldwide consortium defines and agrees on a distributed architecture and basic set-top requirement that presents a universal means of connecting CPE to the access network and enables volume production of STBs ranging from basic to full-featured; or
- A combination of all three, based on defining interfaces, signaling, and control is ideal. This then allows different telcos and countries to choose set-tops or residential gateways and operate their service network from a common back-office platform. Each access vendor would be free to insert its own added-value features and capabilities, supported uniquely from its access platform, as a means of product differentiation, rather than having the CPE device determine differentiation between system solutions.

3. *Considering In-Home Distribution Related to CPE Alternatives*

The "Achilles heel" thus far in the limited commercial deployment of video over copper has been in-home installation. In the early stages, in-home installation took an average of seven hours per home. It has been refined through better procedures, training, and some technology enhancements to now average less than 3.5 hours, allowing an installer to do two service orders per day.

In-home distribution must serve each television and assume one or more personal computers (PCs) as well. The media available in the home consists primarily of existing coax for

CATV distribution and twisted pair Cat. 3 copper for telephone and modem connection. With VDSL terminated at the network interface device (NID) (where lifeline plain old telephone service [POTS] is split off the drop) or at a modem device (in the distributed model), bandwidth and service must be distributed by using the existing distribution media, or else new media must be installed.

In the current residential-gateway model, VDSL is placed on the existing coax, and a coax in-home "network" is used to provide service to each television. Cat. 5 copper is installed from the gateway to the PC, which is usually not located in the same room as the main television, where the gateway resides. This becomes the most problematic installation piece.

The current distributed model assumes in-home distribution uses Ethernet over Cat. 5 copper to both televisions and PCs. This is no less an installation problem, as the majority of homes is not pre-wired with Cat. 5 and have RJ-45 jacks in each room.

Ideally, broadband wireless technology would be developed to allow VDSL over the air with no in-home wiring required. While this is being studied and may become a future possibility, the solution does not exist today. And there will always be an issue with wireless interference (frequency separation) when the application is in a large residential high-rise apartment or condominium complex.

A third alternative is using HPNA 2.0 for both video and data distribution. HPNA 2.0 can work on either coax or Cat. 3 copper and provide up to 30 megabits per second (Mbps) of bandwidth—more than enough for video and data services.

The good news is the availability of several technology and physical media options to deliver video and data services in the home. Before looking more closely at the pros and cons of each solution, it should be remembered that the objectives for deciding the best topology are as follows:

- The solution that best optimizes the amount of time spent installing in the home
- The solution that is best capable of delivering bandwidth and performance for current as well as future services and applications
- The solution that represents the lowest-cost approach to service installation, including electronics and/or peripheral devices

3.1. VDSL over Coax (TV) and Cat. 5 Ethernet to PC

Reuse of coax in the home to serve multiple televisions, as done with a residential gateway, is a fairly effective approach. With the aid of a balun, VDSL terminated at the NID can be bridged directly onto the coax at the side of the home for distribution. This has several limitations, however:

- It assumes the CATV entry and telephone entry are on the same side of the house or within proximity of one another.
- It assumes that the home is pre-wired with coax. This is true for a majority of homes in the United States and Canada, but much less so in Europe and other parts of the world. It may require new coax distribution be installed in a number of international markets.
- It still requires some changes to coax splitters in the home to maintain a quality video stream.

Second, and more costly, is the need to run Cat. 5 from the gateway, where the main television is to the location of the PC for data service. As an option, HPNA 2.0 may be used. The problem still remains that, in many homes, there is no phone jack located where the main television is, so wiring and installation may still be required.

In-home installation intervals (excluding actual VDSL drop provisioning) ranges from two hours for a single-family home with existing coax to three hours with PC connectivity. Condominium and apartment complexes require less time, in the range of 1.5 to 2 hours. If the home does not have coax distribution placed, or new coax needs to be run to a television, installation increases to 4+ hours per home.

3.2. VDSL Termination, Ethernet Distribution

With most STB manufacturers offering a 10BaseT (or 10/100BaseT) interface, video deployment over DSL (mostly ADSL DSL access multiplexer (DSLAM) architectures) uses Ethernet for in-home distribution. DSL is terminated at an intelligent modem and then distributed to STBs around the home using Cat. 5 copper carrying Ethernet. If multiple televisions and/or PCs are connected, a multiport Ethernet bridge is used.

10BaseT Ethernet provides acceptable bandwidth for delivering video to a set-top device. It is limited, however, in providing sufficient bandwidth to a multistream gateway that may also support data service. And it does not support bandwidth required to offer high-definition television (HDTV) service as a future consideration.

The main issue is in-home wiring cost. Except for the newest homes, the vast majority of homes are not pre-wired with Cat. 5, nor are RJ-45 outlet jacks installed in the home. Therefore, complete new in-home wiring must be performed. This, coupled with the modem costs, may make this solution too costly for any commercial mass deployment. A home with more than one television and a PC could take well in excess of four hours for in-home installation, and it is unlikely that this topology could ever migrate to self-install.

3.3. HPNA 2.0 Home Distribution

HPNA 2.0 offers some interesting enhancements over traditional Ethernet. The distinctions are greater bandwidth and the ability to deliver high bandwidth over existing in home telephony pairs (Cat. 3).

Based on the current Broadcom chip package with HPNA 2.0, both 10/100BaseT and HPNA outputs are provided. HPNA 2.0 in its present form can deliver somewhere between 15 and 18 Mbps of bandwidth for home distribution. While this sounds acceptable for multiple video streams as well as data, the reality is that Ethernet-based packets are not ideally suited for constant bit-rate (CBR) video, and mixing bursty data traffic presents a potential problem. If HPNA 2.0 is to be used for the delivery of both video and data inside the home, it probably requires a much larger "pipe" to compensate for the bursty traffic, probably in the range of 30+ Mbps. Technical developments are pushing HPNA 2.0 to a much higher rate, and it could be feasible to deliver a whole home video and data solution over existing telephone wire.

HPNA 2.0 offers the option to deliver service via telephone pairs or coax. Thus, it is flexible in its in-home implementation to best reduce installation time. Putting HPNA 2.0 onto coax requires only an inexpensive balun.

Once the bandwidth is increased to support mixed traffic, HPNA 2.0 will be ideal for delivering both video and data to multiple devices in the home. The trade-off between asynchronous transfer mode (ATM) and Internet protocol (IP), when it comes to mixed broadband transport, is related to the lack of dedicated PVRs for service mixing inherent in ATM. So the only practical solution, as mentioned, is to increase the size of the pipe.

Current-generation chips supporting HPNA 2.0 appear to deliver around 15 Mbps of bandwidth. This may be acceptable for two broadcast streams and data but likely not for three broadcast streams. The next iteration of chips, due in 2002, should deliver 30 Mbps, which would be a large enough bandwidth pipe for three video streams and data. Prototype lab chips indicate the capability to increase bandwidth to more than 100 Mbps. This is a compelling argument for HPNA 2.0, since there appears to be a definitive migration path for higher bandwidth.

One approach would be to start with both a 10/100BaseT and HPNA 2.0 interface. The 10/100BaseT could be used for video, since most DSL STBs today only support a 10BaseT interface. This would require installation of Cat. 5 copper to the STB and does not, again, simplify the in-home installation issue more than the other choices. HPNA 2.0 could be used to deliver data to PCs using existing Cat. 3 telephone lines in the home. It also offers the ability to connect other IP devices, such as security systems, smart appliances, etc., over telephone lines within the home and in a variety of locations (such as the kitchen). The converse can also be deployed, especially in homes where only one to two televisions are connected. In this case, HPNA 2.0 with a balun could be used to serve televisions over coax (or directly over telephone lines if the STB had a HPNA 2.0 interface) and the 10/100 port used for PC connectivity for data.

A major benefit to HPNA 2.0 is that it uses a medium that is familiar to the telephone company and has the flexibility of running over telephone lines or copper. This makes HPNA 2.0 flexible in both the U.S. and international markets, where the availability of coax varies greatly. It also has a bandwidth migration path that forecasts much shorter in-home installation times. Another advantage would be locating the smart modem at the most accessible telephone jack to access home distribution, rather than trying to network access from the main television, where a gateway is located.

4. Final Remarks

This paper has talked about the capital costs for the cable operator being heavily weighted on the fixed infrastructure side, with the CPE and in-home costs being the reverse for the telcos. To put this in perspective, on an equal take-rate basis, the installed costs for two-way digital CATV with two digital STBs per home (the standard service offering of most U.S. cable operators) is more then double the costs for the telco using DSL and digital STBs. The telco problem is more business-case–related, as take rate is a higher risk and the return on investment (ROI) acceptable to a cable operator is significantly lower in the mind of the telcos.

There are a number of DSL access systems that offer scalability and performance for video and offer a means to deploy infrastructure at an acceptable cost. Recent advancements have reduced the costs to build a digital video head end, and lower encoding rates have made bandwidth utilization more efficient. The problem, therefore, comes down to CPE and in-home installation.

In a real sense, set-top manufacturers and other consumer electronics producers that could build low-cost STBs have been reluctant to develop products due to the lack of market. Until the worldwide telcos demonstrate a serious plan to deploy video services, the manufacturers best able to drive CPE costs down will be hesitant to enter the market, as volume drives price in most consumer electronics.

There is encouraging news, though. In the past few months, a growing number of set-top and CPE manufacturers have announced plans to develop either "home gateway" devices or low-cost "distributed" decoder set tops. Of particular interest is the concept of a full- featured STB in the home serving the primary television, with subtending inexpensive, small "decoder" boxes at the other televisions. If the preliminary prices are accurate, this distributed approach can match the cost of a non-distributed, multistream residential-gateway device. This may satisfy the needs of both U.S. RBOCs and worldwide PTTs.

As presented, a number of in-home wiring options are available. Unless a high-bandwidth, high-performance, and low-cost wireless broadband solution is developed for video distribution in the home, there will always be a labor cost associated with home wiring. The near-term goal, through the technology options available, is to simplify and minimize the installation, not eliminate it. Whereas several years ago, a single-family home installation to three televisions and a PC took as long as seven hours, today this number is more like 3 to 3.5 hours with technology improvements. With some of the suggested approaches, such as HPNA 2.0, this can hopefully be reduced to 2 to 2.5 hours. In comparing in-home installation to the cable competitor, there is parity at this level, as they spend an average of 2 to 2.5 hours per home to install two-way digital services to two televisions.

MPLS versus PNNI Control-Plane Architectures in 21st Century Service-Provider Networks

Andrew G. Malis
Chief Technologist
Vivace Networks, Inc.
Chair, Technical Committee, and Vice President, Technology
MPLS Forum

1. Introduction

The purpose of this paper is to provide overviews of the private network-to-network (PNNI) and multiprotocol label switching (MPLS) control-plane architectures, compare their salient features, and demonstrate MPLS' superior applicability to meeting today's and tomorrow's public service-provider control-plane architecture requirements, especially in converged multiservice networks.

2. PNNI

2.1. What is PNNI?

PNNI [1] was created for asynchronous transfer mode (ATM) networking and offers extensive capabilities for connection control among the nodes of ATM networks. ATM was originally created in the late 1980s to be the heart of the envisioned broadband integrated services digital network (B–ISDN), which was at that time intended to completely replace time division multiplexed (TDM) voice and data with hardware-based packet switching. At that time, the only type of packets that could be switched in hardware had to be fixed length, and, in addition, these fixed-length packets (called cells) were felt to provide better quality of service (QoS) capabilities than variable-length packets, which was necessary to replace TDM's ability to carry voice traffic. The use of a fixed cell size has led to a number of consequences, including complicated adaptation functions to carry variable-length Internet protocol (IP) and other packets over ATM, and extremely complex QoS service classes and cell-scheduling algorithms.

Because ATM was intended to replace TDM, it has many characteristics in common, including a requirement that all user cells (carrying both voice and data) be carried over connections that were explicitly established before any cells could be sent. Part of the connection establishment process was to determine the connection's routing through the network, which was then fixed by the connection-establishment process.

PNNI was created by the ATM Forum to provide a standards-based method of routing and signaling connections in ATM networks. Its capabilities include the following:

- *Addressing:* PNNI uses 20-byte network service access point (NSAP) addresses, which were originally invented for International Organization for Standardization (ISO) connectionless network protocol (CLNP).
- *Link-State Routing:* PNNI includes all of the functionality of a link-state routing protocol, such as open shortest path first (OSPF), including a hello exchange, topology generation, and flooding routing updates (known as topology state packets) on the state of all of the network's switches and links to every ATM switch in the network. This routing information allows each ATM switch to determine the optimal route through the network to set up a new connection. That route is then explicitly carried in the connection set-up signaling.
- *Signaling:* PNNI uses the signaling ATM adaptation layer (SAAL) to encapsulate signaling messages that carry the connection set-up requests from switch to switch along the path chosen by the routing function. The signaling messages themselves are an enhanced superset of the International Telecommunication Union–Telecommunication Standardization Sector (ITU–T) Q.2931 connection signaling protocol, which was created for use in the B–ISDN. A specific virtual path identifier (VPI) and virtual circuit identifier (VCI) are used for carrying PNNI signaling information.
- *Multiservice Capabilities:* PNNI is an ATM–only control protocol. To carry other connection-oriented protocols

on PNNI–based networks, additional interworking protocols are required. For example, frame-relay carriage requires the use of the Frame Relay Forum's FRF.5 specification on frame-relay/ATM network interworking, which adds complexity to the network. In addition, carrying frame relay and other packet-oriented networks on a PNNI network requires a one-to-one mapping between frame-relay and ATM connections. (Note that FRF.5 does include an optional many-to-one mapping, which has never been operationally deployed.)

PNNI can be thought of a single monolithic whole encompassing all of the aforementioned functionalities. Its architecture can be visualized in *Figure 1*.

The interfaces on each side are to other switches in a PNNI–based network or across private ATM network boundaries

2.2. PNNI Control-Plane Issues

PNNI has a number of issues that affect its suitability for use as the signaling control plane in a multiservice network. These include the following:

- PNNI was designed for, and is highly optimized for, use on ATM–only networks. It does not natively support signaling or carriage of other protocols—this must be supported by using additional functionality, such as provided by FRF.5 for frame relay.
- Although PNNI was originally intended for use between different vendors' ATM switches, because of its complexity, multivendor interoperability was never really achieved and practically all PNNI networks are single-vendor deployments. In addition, many service providers with ATM networks have opted to use their vendor's proprietary control-plane protocols, such as Lucent/Cascade's Virtual Network Navigator and Marconi/Fore's ForeThought, due to their added functionally and vendor differentiation when compared to PNNI.
- It should be noted that PNNI does not, of itself, provide any particular QoS services for user traffic. Rather, PNNI has the capability of including, in its connection signaling messages, the ATM service class that is requested by the user. It is the job of the ATM switch, where the connection originates, to find a path through the network that can carry the service class and the job of each ATM switch on the path to provide that QoS to the ATM cells on the connection.
- PNNI's use of NSAP–format addresses is problematic, since although this address format was invented for ISO CLNP, these addresses (like CLNP) have never been operationally deployed. It also means that IP deployments on PNNI–based networks need a network translation function.
- PNNI signaling also has a number of scalability concerns. As previously mentioned, carrying frame relay and other packet-oriented networks on a PNNI network requires a one-to-one mapping between frame relay and ATM connections, which will create a large number of ATM connections to be managed and signaled. In addition, because every ATM switch in the path of the connection must actively participate in the PNNI signaling, this means that interior switches in the middle of the network need the signaling and switching resources to support very large numbers of ATM connections, which has a direct effect on the scaling abilities of such networks.

FIGURE 1

PNNI Architectural Reference Model

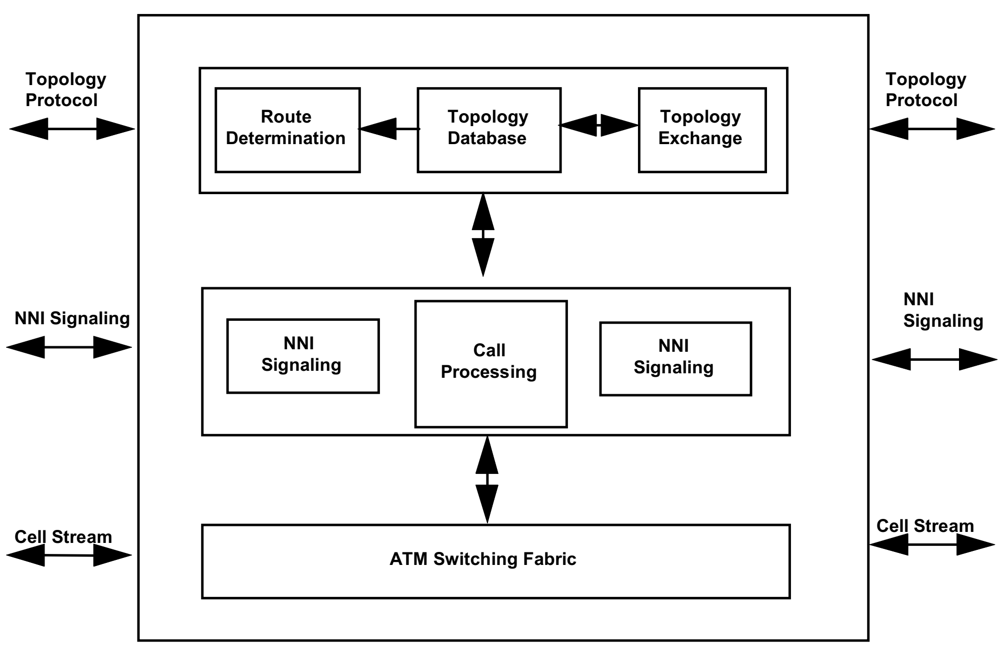

FIGURE 2

IP over ATM Network

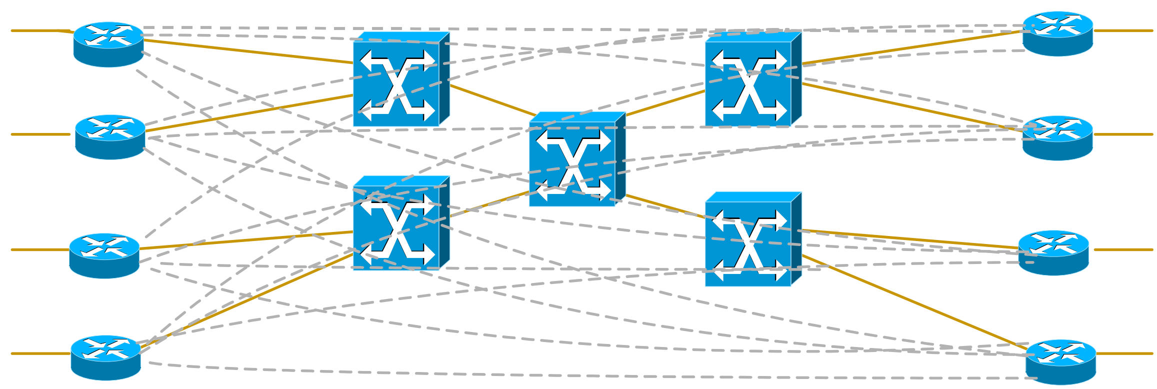

3. MPLS

3.1. What is MPLS?

MPLS [2] is a set of IP–based signaling protocols that were designed by the Internet Engineering Task Force (IETF) to allow the use of connection-oriented traffic engineering in IP networks and to enhance the operation of IP over existing ATM networks. It has lately been enhanced with additional functionality, which will be subsequently discussed further.

Prior to MPLS, router-based IP networks used hop-by-hop routing of IP packets, which tended to send IP packets over a small number of possible paths though the network, causing congestion on some paths and underutilization on others. During the relatively short period when ATM switches tended to be faster then IP routers, IP backbones were built around a core of ATM switches. The ATM switches allowed the use of ATM connections to perform traffic engineering in order to evenly spread the load of IP packets on the network's links, but with a number of costs. The process of adapting variable-length IP packets to ATM's fixed-length cell produced a large amount of overhead, known as the *cell tax*, which could be up to 30 percent of a link's bandwidth. Typically, network operators used a full mesh of ATM connections to interconnect the IP routers, which used a large number of connections and other resources in the ATM network. In addition, an Internet service provider (ISP) was required to procure and operate two sets of network infrastructure equipment—both ATM switches and IP routers. This is illustrated in *Figure 2*.

In *Figure 2*, the solid lines are physical connections and the dashed lines represent a full mesh of ATM VCs that are being used to interconnect the routers around the core of ATM switches.

Once routers began to use hardware switching in application-specific integrated circuits (ASICs) to match and even surpass ATM's switching speeds, the only use for ATM switches in ISP backbones was to provide traffic engineering. To reproduce this functionality in routers, the concept of label switching for IP packets was invented in late 1996. This simply prepends variable-length IP packets with a fixed-length 20-bit label to quickly switch the packets through the routers along a path that was independent of its destination IP address, thus allowing traffic engineering in routers without the use of ATM switches. As a result of its genesis in the IP routing community, MPLS is much more "open" than PNNI, and it has reaped the benefits of more than 30 years of packet switching technology experience.

3.2. MPLS Advantages

The various aforementioned aspects of PNNI can be compared to their MPLS counterparts:

- *Addressing:* MPLS uses the same 32-bit addresses as IP. In the future, MPLS can be extended to IPv6, just as IP itself is being extended.
- *Link-State Routing:* MPLS uses the existing IP link-state routing protocols, OSPF and intermediate system–to–intermediate system (IS–IS), with a small set of extensions to provide additional information for traffic-engineering purposes. A service provider has the flexibility to choose one or the other routing protocol; neither is mandated by MPLS.
- *Signaling:* MPLS has defined three signaling protocols, again with the network provider making the choice as to which to use in their network. The three signaling protocols are RVSP–TE (traffic-engineering extensions to the existing resource reservation protocol [RSVP]), LDP (label distribution protocol, used for running MPLS in a basic mode of operation), and CR–LDP (constrained routing extensions to LDP, also used for traffic engineering). All MPLS implementations contain LDP and either one or both of RSVP–TE and CR–LDP.
- *Multiservice Capabilities and Scalability:* The MPLS architecture includes the capability to set up multiservice connections to carry ATM, frame relay, Ethernet, TDM, and other Layer-1 and Layer-2 services. When MPLS–based networks are used to carry these multiservice connections, such as ATM or frame relay, many multiservice connections are hierarchically carried in a single MPLS LSP. In addition, the ATM and frame-relay connections are only visible to the endpoint switches where they terminate; the switches in the interior of the network do not participate in the signaling. This hugely increases the scalability of

MPLS networks for multiservice traffic when compared to PNNI.

- *QoS:* MPLS provides the ability to allocate bandwidth to LSPs and to associate LSPs with QoS service classes to ensure that user applications will receive the required end-top-end QoS from the network. In addition to MPLS' standardized QoS capabilities, Vivace Networks has incorporated the lessons learned from ATM switch design to provide its customers with the best of both worlds: MPLS' signaling scalability and packet-based forwarding combined with the capabilities to provide rigid QoS guarantees for its customers.

3.3. MPLS Flexibility and Future Direction

In comparison to PNNI's monolithic design of one protocol handing topology distribution, routing, and connection setup, MPLS was designed in a modular manner. MPLS has been built upon already existing and proven protocols in use in packet-based networks, such as OSPF and IS–IS, for topology distribution and routing. In addition, MPLS allows service providers the specific choice of which routing and signaling protocols they wish to use to best suit their requirements, while PNNI attempts to cover all possible operational situations, with no service-provider flexibility allowed. All of this makes MPLS the perfect control-plane architecture for use in packet-based multiservice backbone networks.

In comparison to MPLS, PNNI represents a technological dead end. The ATM Forum, where PNNI was standardized, has vastly shrunk in size and interest, and it is extremely doubtful that there will ever be another revision to PNNI. In direct contrast, the IETF and the Optical Internetworking Forum (OIF) are cooperatively working on extending MPLS' capabilities to become, as generalized MPLS (GMPLS), the control plane of choice for service-provider core optical networks. In addition, the MPLS Forum is complementing the IETF's MPLS standardization work by actively extending MPLS' application space and standardizing MPLS' protocol conformance and functional testing.

4. Summary

Table 1 summarizes the salient points of PNNI and MPLS.

5. References

[1] The ATM Forum, *Private Network – Network Interface Specification Version 1.0 (PNNI 1.0)* (March 1996), af-pnni-0055.000.

[2] Rosen, E., et al., *Multiprotocol Label Switching Architecture* (January 2001) IETF RFC 3031.

[3] Davie, B., et al., *MPLS using LDP and ATM VC Switching* (January 2001) IETF RFC 3035.

TABLE 1

PNNI and MPLS Summary

Functionality	PNNI	MPLS	Comments
Data units	Fixed-length cells	Variable-length packets	MPLS is much better suited to today's IP–based protocols.
Standardization background	ITU–T and ATM Forum, based on B–ISDN	IETF, based on IP routing and signaling	PNNI is one monolithic whole and difficult to adapt to changing times, while MPLS is modular in design, and each aspect can be updated on its own.
Typical data rates	T1/E1 to 2.5 Gbps	9.6 kbps to 10 Gbps, 40 Gbps in development	
Addressing	20-byte NSAP addresses	32-bit IPv4 or 128-bit IPv6 addresses	ATM's different addressing structure requires the use of a mapping function when using IP over ATM.
Signaling	Superset of ITU–T Q.2931	LDP, CR–LDP, RVSP–TE	
Routing	Included as a part of the protocol	OSPF or IS–IS	MPLS makes use of routing protocols that were designed specifically for IP.
Network scalability	Two-level hierarchy (VCI and VPI)	Infinite-level label stack	ATM VCI and VPI fields are limited in size and make network scaling difficult.
QoS	Not provided by PNNI per se, but by the underlying ATM switches	Included in the MPLS architecture	In addition, Vivace Networks has innovated to provide QoS capabilities over and above those in the base MPLS architecture.

TABLE 1

PNNI and MPLS Summary (cont'd)

Protection and resilience	PNNI can take multiple seconds to re-establish connections that break due to a switch or link failure	MPLS has 50 ms LSP restoration capabilities	PNNI relies on underlying SONET/APS to provide 50 ms restoration; MPLS has it built in.
Connection set-up resilience	PNNI includes crankback to try alternate routes if a problem occurs during connection set-up	MPLS includes crankback to try alternate routes if a problem occurs during connection set-up	PNNI and MPLS are identical in this regard.
Overhead	ATM cell tax (up to 30 percent for packet-based services)	Label stack in front of packets (32 bits per label)	MPLS is much more efficient for packet-based services.
Privacy and security	ATM uses VC separation to prevent mis-delivery of user data	MPLS uses LSP separation to prevent mis-delivery of user data	
Network backbone convergence	As discussed, PNNI and ATM have serious scaling and efficiency issues when used as the basis for multiservice converged backbone networks.	MPLS is able to efficiently and scalably concentrate the traffic from Layer-1, -2, and -3 services at the edge and move them transparently across the core.	
Operational Cost	IP over ATM requires two separate control-plane architectures.	Single converged backbone	Converged MPLS core simplifies operation and maintenance and reduces the troubleshooting time under network outages.
Multivendor interoperability	Very few multivendor deployments	Proven in real network deployments	
Migration to new optical core technology	PNNI is incompatible with the forthcoming OIF optical interfaces.	The OIF UNI and NNI interfaces are based on GMPLS, a superset of MPLS.	MPLS eases the migration to the forthcoming GMPLS–based optical core and enables the direction cooperation of service provider data and optical switching equipment.

Data-Centric Design and Its Impact on Storage and Network Architecture

Rao Mikkilineni
Chief Executive Officer and President
enSANET Inc.

Andy Helland
Director, Product Management
LightSand Communications, Inc.

Abstract

When computers were making a major impact on enterprise productivity in the 1980s, the world was starved for processing power. The most important feature in a computer system was how fast the central processing unit (CPU) ran. We were, in effect, CPU bound. As the world moved into the decade of the 1990s, CPU power became more and more plentiful. However, the world soon perceived an ever-increasing need to network computers together in a globally scalable fashion. We shifted our attention to our new weak spot: networking. Now that processing power and connectivity abound in the first decade of the new millennium, we have shifted our attention to storage. Since our attention is now on data and how to make sure that it is safe and available in the location in which it is required, several new requirements are being levied on the world's infrastructure. These are predictable quality of service (QoS) as well as a security infrastructure to guarantee that data is protected and authentic. These new requirements will propel the design of high-bandwidth, low-latency infrastructure for the world with end-to-end protection of the data.

Introduction

The decade of the 1980s truly belonged to the CPU. Since Moore's observation[1] that the number of transistors per square inch on integrated circuits has doubled every year, the pace has settled down and continued to double every 18 months. The decade of the 1990s has seen an even faster rate of growth in the network bandwidth capacity that doubled every 9 months. With the advent of the Internet and faster global communications networking, storage requirements are exploding at an even more rapid rate. E-mail, e-commerce data, and continued conversion of documents, pictures, and video to digital storage are just examples contributing the exponential growth rate that is doubling storage every six months. In addition, business focus is shifting to 24x7x365 data availability. Down time has been estimated by the Gartner Group[2] as costing anywhere from $300,000 per hour for industries such as health care and media to $3 million per hour for highly automated businesses, such as the financial and telecommunications industries. More recently, the World Trade Center disaster has highlighted the importance of business continuance using back up, mirroring, and other reliable disaster recovery mechanisms. The first decade of the 21st century may well belong to storage infrastructure innovation.

Historically, storage has stayed very close to the computer, where data reads and writes originate. This has allowed the storage industry to develop very high throughput and low-latency infrastructure to optimize the storage access. However, as bandwidth started to explode, the economics has shifted to networked storage. Networked storage allows the sharing of storage resources across the enterprise and reduces operation and management costs. Network-attached storage, Fibre Channel storage-area networks (SANs), iSCSI, and InfiniBand technologies are examples of such networked storage solutions.

As networked storage becomes ubiquitous, two new issues need attention. First is the latency introduced by increasing the distance between the server and storage through the network, and second is the security that must be addressed as mission-critical data starts to share networking infrastructure in both enterprise and public networks. Storage transactions pay a premium for accuracy of the data and require acknowledgement when a write is accomplished before the transactions are committed. As the networks increase the distance between the server and the storage device, management of both latency and security becomes critical. In this paper, we examine the impact of latency and security on storage networking and identify key requirements that must be satisfied by the networking infrastructure. Storage caching among distributed data centers will become an important way to not only utilize the efficiencies provided by high-bandwidth pipes at ever decreasing costs, but also to mitigate the limitations imposed by latency over distances. As storage comes out of the traditional data-center closet and becomes globally networked, many of the security risks will be addressed with similar technologies that have provided solutions for the traditional Internet protocol (IP) networks.

Emergence of Storage Networks

First-generation storage devices were directly attached to their servers. Large computer systems from IBM used the enterprise system connection (ESCON) interface, while small systems use the small system computer interface (SCSI). The SCSI protocol was developed in 1981 to interconnect disk drives and other peripherals to small computers. It was submitted to the American National standards Institute (ANSI) standard committee and was formally adopted as a standard in 1986. The initial version of the SCSI interface allowed up to seven devices to be attached to the computer. Since SCSI was intended to interconnect devices within the computer itself, the interface was left as parallel and large ribbon cables were used to connect peripherals. Various extensions to SCSI increased the speed of the interface as well as the number of users that could simultaneously use the bus. The latest version of SCSI allows up to 15 devices to be attached to the server and supports data rates up to 320 megabits per second (Mbps).

As more complex storage arrays were assembled, two attributes of SCSI drove the storage industry to develop Fibre Channel SANs. These were the distance limitation of SCSI (25 meters) and the lack of addressable devices (15). To mitigate these limitations, Fibre Channel networks were adopted using the Fibre Channel protocol to provide a switched fabric for the interconnection of multiple SCSI devices. Fibre Channel provides a switched 100 Mbps interconnect between devices. New standards have doubled this data rate to 200 Mbps equipment, with this equipment available now. Not only did Fibre Channel allow for hundreds or even thousands of devices to be interconnected, but it also increased the maximum distance from 25 meters to 10 kilometers.

The Fibre Channel standard allowed SANs to be built that provided mirroring, redundant architectures, and scalability within the data center. Fibre Channel provided lots of bandwidth and very low latency in interconnecting storage systems. This architecture works well for many applications but falls short of our ultimate goal of providing continuously available data. When considered as one large system, the entire data center is vulnerable to natural disaster. For this and other reasons, we need to distribute our data geographically. The next section will discuss this need in greater detail.

The Need for Distributing Data

Since our new emphasis in the 21st century is on ensuring that data is safe and available in the location(s) in which it is required, we focus on two main reasons for data distribution: business continuance and content caching. The first reason covers all aspects of protecting the data of the enterprise and protecting the systems used to deliver data to the end consumer. The second reason addresses the economics and performance of caching data in disk systems local to the end user versus having all end users access the data from a centralized data center. This provides wideband access to data for all users by eliminating the latency and cost of pulling all data through the long pipes from a centralized source. We will address these topics one at a time.

Business Continuance

Quite simply, business continuance encompasses all aspects of data storage and system design that impact the safe storage and reliable delivery of data. Storage-system designers recognized long ago that a disk-drive failure constitutes a catastrophic loss of data. In response to that danger, arrays of two or more disk drives were deployed in which the data is written in multiple devices. If one drive failed, its data could still be found in one or more other drives. These mirrored storage arrays were typically located in the same facility as a convenience to the staff that managed the storage system. Soon, however, it was recognized that a catastrophic natural disaster could also destroy all of the data. Even though the data was mirrored within the data center, a flood or earthquake could still destroy all of the data in the entire facility.

To mitigate the threat of a natural disaster, storage-system designers created mirrored data centers that stored a copy of the data in a location that was physically separate from the primary data center. Thus, if an earthquake or other natural disaster destroyed the primary data center, the data would still be safe in the mirrored site. To date, these systems have been extremely expensive given the need to interconnect the sites with high bandwidth and low latency. They have typically been used by only the very high-end consumers of data who can measure the loss of data or the loss of service in thousands or millions of dollars per hour.

As the world has achieved better and more affordable network connectivity, more and more users are beginning to protect their data in this fashion. Some are deploying Fibre Channel over a higher-level protocol such as IP in an effort to take advantage of the prevalence of IP networking infrastructure. Others, seeking a higher level of performance, are deploying Fibre Channel directly onto low-level transport mechanisms such as synchronous optical network (SONET) or dense wavelength division multiplexing (DWDM).

Regardless of which method of networking is used, the interconnection of storage between remote sites will continue the pressure on carriers to provide a higher level of QoS.

Content Caching

We expect content caching to emerge as a new paradigm for storage networks. There are two reasons for this: performance and economics. Both reasons ultimately derive from the increased latency that results from locating data centers farther and farther apart from each other.

Let's consider the model of a Web hosting service. If they operated one data center (say in San Jose), then all Web pages served by them would be generated in San Jose and served to clients across the country or even the world. If the Web site contained a significant amount of data (for example, high-resolution photographs), then the server would have to push those large data files to every user (regardless of their location). As the number of users increases and as the distance between the server and the user increases, the distance-bandwidth product of the data in motion will increase dramatically. Ultimately, this drives the cost of providing data to the user, since it is a function of both bandwidth and distance. To complicate matters, the requirements for bandwidth will continue to increase as consumers seek ever-richer forms of media entertainment (streaming video, etc.).

Since we expect the demand for bandwidth to grow considerably, all that is left to us is to reduce the distance over which the data is served. In addition to decreasing the cost of the network infrastructure (distance-bandwidth), this has the added benefit of decreasing the time necessary for the server to respond to the user (less flight time for the packets in motion). This yields a more responsive system with greater throughput. All of these same issues will apply, whether the company is serving Web pages or just managing the distribution of wideband data. The Web-hosting company Akamai has developed an architecture that is based on these principles. Start-up companies such as Accellion are also developing architectures such as this to provide wideband, distributed data access at the edge of the network.

To implement a system such as this, the typical enterprise would need to deploy multiple data centers with identical (synchronized) content. Each data center would serve the region of the country in which it resides. It would provide wideband access to data via a much shorter link than had been used in the past. However, in implementing an architecture such as this, the storage system is faced with the challenge of maintaining synchronized content at different locations. By design, this represents less demand on the network than the centralized model, but nonetheless; it creates an increased need for block-level storage data to be moved across the network.

As we discussed, any time that block-level transport of data is involved, high bandwidth and low latency are crucial to maintaining high performance.

The Implications of Data-Centric Design on Security Infrastructure

Because storage was confined close to the computers where the data originated, storage has traditionally relied on physical security. As the networked storage emerges, this will no longer suffice.

There are three major types of storage networks that are evolving. First, the emergence of Fibre Channel–based SANs that provide low-latency, high-speed data transport between servers and the disk arrays located several miles away from each other using optical transport. Second is the iSCSI technology that allows carrying SCSI commands from the server to the disk arrays over traditional IP networks. Third is the future InfiniBand technology that is aimed at replacing current peripheral component interconnect (PCI) buses in the servers with high-throughput, low-latency switched networks. All three types of networks allow the storage to be extended beyond the data center through appropriate bridging with metropolitan-area networks (MANs) and wide-area networks (WANs).

As data makes its way out of the data center into public and private networks, security concerns become important. Storage networks will be subject to the same three main risks associated with other network-connected systems. The primary risk is that of a user gaining inadvertent access to unauthorized data. Another significant risk is an attacker who hacks the storage network or, more likely, one of its attached host servers and obtains administrator credentials. The third risk, applicable in the case iSCCI networks, is a denial-of-service attack that floods the storage network with more requests than it can handle.

First, let us look at the model of the potential threat that storage faces.

Since Fibre Channel is the most prevalent protocol used in the design of SANs today, we will concentrate on the vulnerabilities of Fibre Channel–based SANs. As Fibre Channel is transported over dedicated fiber-optic cables, it will remain relatively secure even as it travels out of the data center. The fiber-optic cables are not readily accessible, and it would be difficult, if not impossible, to perform passive monitoring without alarming the legitimate users of the fiber-optic cable. However, attempts extend the distance by combing Fibre Channel and IP protocols using Fibre Channel over IP (FCIP) expose the data to all of the same threats and attacks that have plagued IP for years.

As part of the FCIP request for comment (RFC), IPSEC is being added to protect the underlying Fibre Channel frames. This should provide cryptographic integrity as well as device authentication. Lately, proposals have started to emerge for end-to-end Fibre Channel security mechanisms (FCSEC). These would provide cryptographic data integrity and device authentication regardless of the physical transport used for the Fibre Channel data.

Security capabilities are also being incorporated into the storage fabric itself through logical unit number (LUN) masking and zoning. LUN masking creates subsets of storage within the SAN virtual pool and allows only designated servers to access the storage subsets. LUN masking provides the ability to create and manage virtual pools of storage, a capability essential for storage service providers (SSPs) to allow multiple customers to manage their own storage without the security risks. Fibre Channel zoning achieves a similar result at the switch by limiting access from a given device to only certain other devices in the fabric. To date, the zoning techniques that have been deployed have used soft zoning. In this model, a user device never discovers the existence of devices to which it should not talk. The system then relies on the integrity of the initiator to not seek to access other devices. With soft zoning, packets that are sent from a given device to an unauthorized device will still pass through the fabric. Emerging proposals for hard zoning will add packet filters to the switch fabric itself so that communication between unauthorized users can never occur.

With iSCSI protocol, data is carried over IP networks, and hence all the security concerns of the public and private networks will apply to storage over IP. These same concerns also apply to InfiniBand when extended over public and private networks.

The Implications of Data-Centric Design on Network Infrastructure

As the philosophy of data-centric design catches on, there will be a renewed emphasis on the availability of a high-performance end-to-end network infrastructure. As we have discussed, storage systems are particularly sensitive to maintaining high-bandwidth and low-latency interconnection. This can become a significant issue when storage networks are faced with the fundamentally different philosophies that have driven the architecture of the Internet as we know it today.

Today's Internet

Today's Internet was designed with scalability as the prime requirement. In this sense, guaranteed performance (even at a low level of QoS) was sacrificed so that the architecture could scale to enormous size. As of October 2001, almost 130 million hosts had been registered in the public Internet (see www.netsizer.com). The scalable architecture has clearly worked, but it has done so at the expense of guaranteed performance. The core element of today's Internet is the router, designed for the flexible, connectionless transport of IP packets. Through the clever design of hierarchical routing protocols, the router has proven its ability to interconnect millions of hosts. It accepts any and all packets regardless of its load and routes them to the appropriate port so that they can work their way to their final destination, one hop at a time.

In the event that a router reaches capacity, it has a very simple algorithm for managing that excessive load: It drops packets. It does so without any guilt or remorse, since every expectation of IP routed transport has been set to "unreliable." To compensate for the unreliable nature of IP transport, retransmission protocols such as transmission control protocol (TCP) are used to sequence and error-check data. Dropped or damaged packets are identified and then requested for retransmission. Unfortunately, the sensing process and the retransmission process take time. This can significantly increase the latency of the transmission. Further, the TCP algorithm will throttle back its transmission window, effectively lowering the transmission bandwidth as well. Since most applications that run over the Internet are asynchronous, they have do not have rigid requirements for latency or bandwidth. The TCP/IP protocol suite has repeatedly demonstrated its ability to provide scalability. Unfortunately, it is difficult to guarantee any set level of performance using TCP/IP over a lossy, routed infrastructure.

Tomorrow's (Storage Friendly) Internet

The interconnection of storage systems for business continuance or content caching may well turn out to be the "killer app" that will drive the Internet toward predictable, high performance.

In the near term, storage-system designers will compensate for the lack of guaranteed quality in public infrastructure by working around the data network. For example, storage networks will be extended using various technologies that allow direct transport of Fibre Channel over asynchronous transfer mode (ATM), SONET, or DWDM. These will mostly be deployed in point-to-point networks between mirrored data centers.

Many new data networking models are being tried. All seek to provide the "holy grail" of infrastructure: predictable, high performance. Among the models being tried are ATM (rapidly falling from favor), multiprotcol label switching (MPLS) (still on probation), and switched metro Gigabit Ethernet. All of these avoid the "drop when congested" behavior of the router. All use some form of flow control or flow-rate management to ensure that excessive data is not transmitted over the network.

Figure 1 shows the SAN interconnection through various options.

It is not clear which of these will emerge as the ubiquitous winner. However, it is clear that there is a tremendous need for predictable QoS in the next generation of data networks. Storage networking is at the front of the line demanding this quality.

Conclusion

The emerging "real-time enterprise" has created a demand for business that operates in a non-stop, on-demand fashion. This has created a need for business continuance solutions that also operate in a non-stop fashion. To meet this need, mirrored data centers will become more common. In addition, storage will become a truly networked asset that can be secured and managed. Networked, geographically dispersed storage, with synchronized content, will provide higher performance to end-users.

However, since storage systems have traditionally enjoyed very low latency and high throughput, there will be new demands placed on the new world of network interconnect to minimize the latency. There are two ways in which the

FIGURE 1

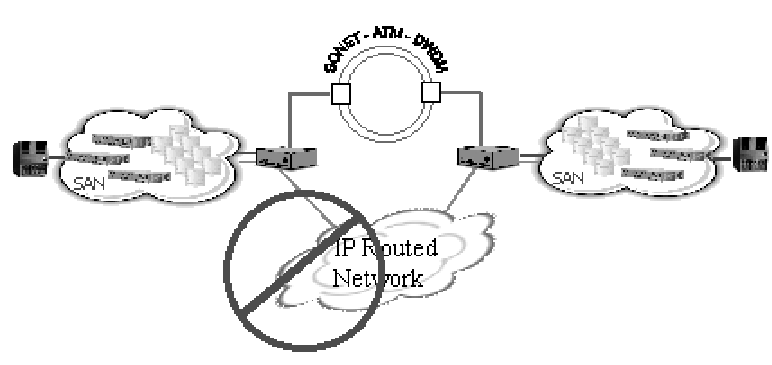

latency requirement will be addressed. Initially, there will be bypass networks deployed that provide point-to-point high-bandwidth, low-latency interconnection. These will be based on SONET or DWDM technologies. Later, interconnect technologies—such as ATM, MPLS, or more likely switched Gigabit Ethernet—will emerge to provide a more flexible interconnection without sacrificing performance. As these technologies stabilize within the world's network infrastructure and as their performance increases, storage systems will move to adopt them because of their ubiquity and lower total cost of operation (TCO).

Notes

1. Gordon E. Moore, "Cramming More Components onto Integrated Circuits," *Electronics*, Vol. 38, No. 8, April 19, 1965.
2. Gartner Group report, 2000.

Softswitch Is a Disruptive Technology

Frank Ohrtman
Telecommunications Consultant and Degree Candidate,
Interdisciplinary Telecommunications Program
University of Colorado

Introduction

The first deployment of the No. 4 electronic switching system (ESS) Class-4 switch[1] in the North American market was in 1976[2]. The performance and operational savings of the No. 4 ESS led to the replacement of all of the preceding technology, namely the No. 4 crossbar switch, within about a decade[3]. The replacement of the No. 4 crossbar switch by the No. 4 ESS is an example of a sustaining technology. A new technology known as Class-4 replacement softswitch is now entering the market and has the potential to displace the Class-4 switch in certain markets.

Definitions

In the mainstream circuit-switched telecommunications architecture, the Class-4 switch interconnects Class-5 switches. Class-4 switches provide traditional toll (long-distance) services.

A new technology has come on the market known as the Class-4 replacement softswitch. Softswitch seeks to surpass the functionality and performance of legacy circuit-switched Class-4 switches by providing lower cost (both in purchase and operations), smaller form factor, distributed architecture, and open standards (see *Table 1*).

In his 2000 business book, *The Innovator's Dilemma*, author Clayton Christensen describes how disruptive technologies have precipitated the failure of leading products and their associated and well-managed firms. Christensen defines criteria to identify disruptive technologies regardless of their market. These technologies have the potential to replace mainstream technologies and their associated products. Disruptive technologies, as defined by Christensen, are "typically cheaper, simpler, smaller, and, frequently, more convenient" than their mainstream counterparts. Furthermore, disruptive technologies "bring to market a very different value proposition than had been available previously ... disruptive technologies under perform established products in mainstream markets."[4]

What is a softswitch? As in the early days of any emerging technology, there is much confusion, if not conflict, over terminology. Currently, there are two popular uses of the term "softswitch." According to the International Softswitch Consortium (ISC), a softswitch (a.k.a. call agent, call server, or media gateway controller), is a device that provides, at a minimum, the intelligence to control connection services for a media gateway, the ability to select processes that can be applied to a call, routing for a call within the network based on signaling and customer database information, the ability to transfer control of the call to another network element, and interfacing to and support of management functions such as provisioning, fault, billing, etc.[5]

A softswitch *solution* is defined by its components. A softswitch solution is comprised of a media gateway controller (also known as a gatekeeper), a voice over Internet protocol (VoIP) gateway (also known as a media gateway), a signaling system 7 (SS7) signaling gateway, and an applications server. This paper will use the term "softswitch" to denote a softswitch solution.[6]

Markets

Over the past few years, dominant North American interexchange carriers (IXCs) have seen their respective market-shares erode. One reason for this is the lowering of the technology barrier to entry into their market. Previously, a long-distance service provider would have to invest great sums in Class-4 switches and nation-wide circuit-switched networks. The high cost of that technology reserved this market for only the very well capitalized. This lowering of the barrier to entry is made possible by new softswitch technology that is cheaper, smaller, and offers distributed architecture and open standards.

Service providers speak of a telecommunications market where voice and data are converged and the subscriber enjoys highly efficient Internet protocol (IP) services "desktop to desktop." This is called a *converged* network (see *Figure 2*). The vast majority of the Class-4 switch market was designed and installed before the advent of converged or IP networks, where voice and data were handled via separate channels. These are referred to as *legacy* networks (see *Figure 1*). For the purposes of this thesis, the author refers to the transition network as a *converging* network (see *Figure 1*). To define the markets for Class 4 and softswitch, it is important to understand that legacy markets applies to legacy networks, where voice and data are separate networks; that a converging market applies to converging networks, where, in most instances, the legacy infrastructure remains of Class-4 and Class-5 switches at the periphery of the network while the core of the network

TABLE 1

Comparison of Class 4 and Softswitch

Class 4	Softswitch
Reliable (five nines)	Not as reliable as Class 4
Scalable (more than 100,000 ports)	Not as easily scalable as Class 4
Good quality of service (QoS)	QoS inferior to Class 4
Established features and applications	Features and applications not defined—possibly greater
Circuit-switched	Packet-switched
Expensive	Less expensive than Class 4
Centralized architecture	Distributed architecture
Large form factor	Small form factor
Proprietary standards	Open standards

is IP providing efficient long-distance voice transport (see *Figure 1*); and that a converged market applies to a converged network where voice and data are handled on one network (see *Figure 2*).

Softswitch is best suited for the converged market, where the majority of long-distance voice traffic is transported on converged IP networks and circuit switching is not present. Mainstream Class-4 technology was designed for legacy markets and is largely unsuitable for the efficiencies demanded in the converged market. Softswitch is gaining market-share at the expense of Class-4 in the converging market.

FIGURE 1

Legacy and Converging Networks with Class 4 and Softswitch

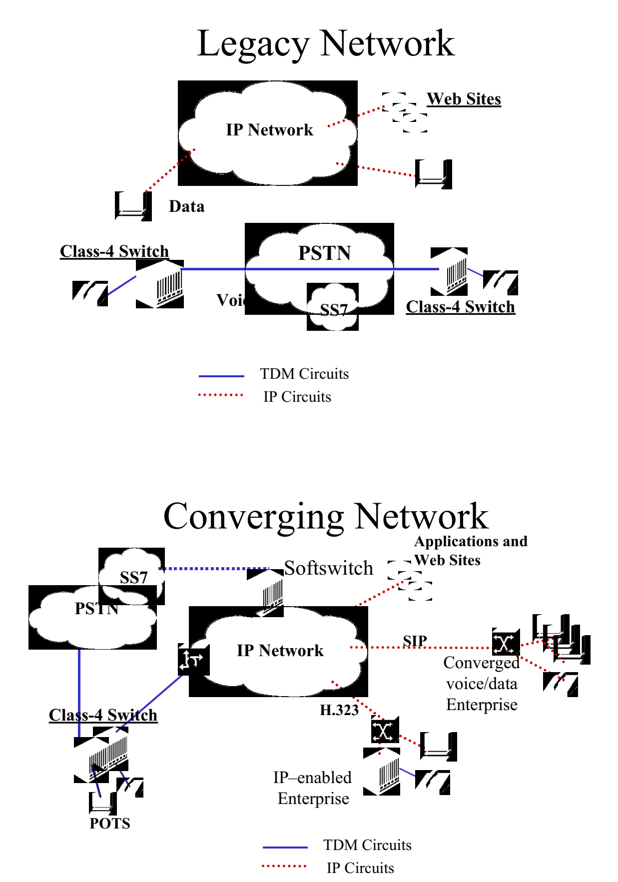

Advantages of a Class 4

Successful companies provide their customers with what their customers demand. In the author's experience in selling switching equipment, service providers demand reliability, scalability, QoS, and well-established features in their phone switches. Class-4 vendors have hybridized the product specifically for the North American legacy market. Its advantages are reliability, scalability, QoS, and well-established features.

Reliability

The focus of debate on this issue in the industry revolves around the question of how many "9s," or how much reliability, a product can deliver. The public switched telephone network (PSTN), or legacy voice network, boasts "five 9s," that is, 99.999 percent reliability. This equates to 5.25 minutes of downtime per year.

Scalability

For a variety of reasons, a Class-4 switch must scale to the maximum possible at the upper end. Given the centralized architecture of the PSTN, "bigger is better." There are approximately 1,400 Class-4 switching offices in the United States. A switch with 100,000 ports delivers advantageous economies of scale. Nortel's DMS-250 Class-4 switch scales up to approximately 100,000 ports. The other matrix for scalability is the call processing power of the switch. The DMS-250, for example, can process 700,000 busy-hour call attempts (BHCA).[7]

Quality of Service

The voice quality of the North American PSTN is the standard for telephone service worldwide. The Class-4 switch was the first deployment of time division multiplexing (TDM) utilizing a 64 kilobit per second (kbps) circuit, which remains the standard to this day. Service on the PSTN is known for the absence of echo, crosstalk, latency, dropped or blocked calls, noise, or any other degradations of voice

FIGURE 2

The Converged Network
Note the absence of Class-4 and Class-5 switches.

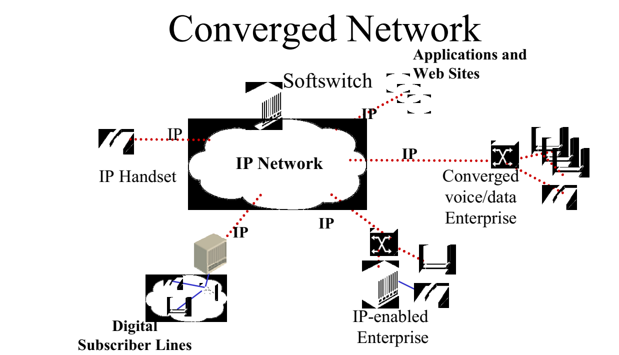

for most long-distance service providers. Any technology that would challenge its dominance must match Class 4 in performance and offer advantages of its own.

Advantages of Softswitch

Softswitch Is Cheaper Than Class 4
Softswitch is cheaper than Class 4 in purchase cost and operations, administration, maintenance, and pProvisioning (OAM&P). A recent study by a major carrier found that packet equipment was 70 percent less expensive than traditional voice equipment, and data access lines were 60 percent to 80 percent cheaper than voice lines. The maintenance of packet networks was 50 percent less expensive, while provisioning was 72 percent lower.[9]

It is estimated that a service provider's largest cost is the ongoing expense of running the network and the switches themselves. Networks must be deployed, maintained, repaired, monitored, and upgraded. Aberdeen Group found that competitive carriers spend 60 to 70 percent of their overall expenses on OAM&P. The ability to manage the entire network via a softswitch solution (including adding and altering specific customer services, upgrading capacity, and fixing faults) from a centralized location leads to considerable time and cost savings.[10] Packet technology–based long-distance carriers maintain that their service-provider customers will be able to offer voice services at 20 percent off traditional rates and claim the underlying infrastructure can be 40 percent less expensive to operate.[11]

quality. Mainstream service providers in North America are unlikely to deploy equipment that offers voice quality at a lesser standard than the Class-4 switch.

The long-distance service providers have owned and operated their own proprietary networks over which they have had total control.

Class-4 Features
With almost three decades of development, the Class 4 has a well-established list of features. Custom local-area signaling services (CLASS) features are basic services available in each local access and transport area (LATA).[8]

DMS-250 system long-distance features can be grouped under two major portfolios: basic and enhanced services. The basic services include such features as 1+, 800/900 service, travel cards, account codes, personal identification numbers (PINs), operator access, speed dialing, and automatic number identification (ANI) screening. Most of the aforementioned features have been standard on the DMS-250 and other Class-4 switches for many years.

In summary, the Class 4's reputation for reliability, scalability, QoS, and features has made it the preferred technology

Smaller Discrete Form Factor
Softswitch is smaller than Class 4. This is an advantage for softswitch. Due to the smaller "footprint," or discrete form factor (smaller size and shape), softswitch power consumption and cooling requirements are also less than the legacy switches due to their smaller size. Smaller hardware size also translates into lower "real estate" expenses. A Class-4 switch is usually located in a central office (CO), where space is a precious commodity and the cost of rack space is high. Softswitch components need not reside in a CO (see *Figure 3*). This opens many possibilities for locating the equipment in less expensive "co-loc hotels," that is, commercial co-location spaces located, in many cases, near COs.

FIGURE 3

Softswitch Solution in the Converging Market

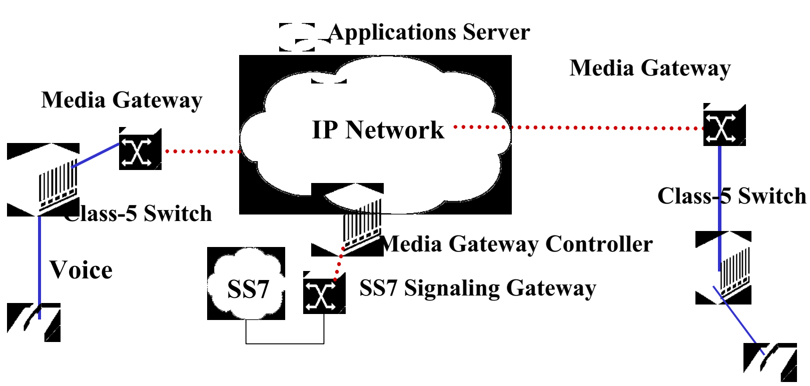

Distributed architecture saves a carrier money in that its components need only reside somewhere on the IP network and not in a CO. Smaller softswitch components can be distributed around the network at less expense than large centralized Class-4 switches.

The smallest configuration for DMS-250 (Nortel's Class 4) is 480 digital signal (DS)–0s (20 T1s). For many locations and service providers, this presents a potentially expensive overabundance of capacity. Many media gateways scale by the T1 or E1 card allowing a greater flexibility in scaling down to capture certain markets and to take advantage of distributed architecture. If a service provider purchases and installs only what they need to get started in a given market and adds capacity only as it is absolutely necessary, then capital is not invested in equipment that cannot immediately generate revenue. The ability of a softswitch to scale *down* is an advantage for many carriers.

Advantages of Distributed Architecture

Long-distance bypass carriers were enabled by smaller gateways to roll out service in diverse markets and took market-share from large mainstream long-distance providers. Assume that service providers that were not traditionally voice service providers—that is, cable-television (CATV) operators, Internet service providers (ISPs), power companies, or digital subscriber line (DSL) providers—wanted to originate and terminate long-distance traffic for their existing customers. Following the long-distance bypass model, they could install media gateways in a number of points of presence (POPs) in tier 1 and 2 cities, and they could potentially take market-share from the legacy long-distance Class-4 carriers. This scenario points to the displacement of Class 4 by softswitch, a disruptive technology (see *Figure 4*).

Open Standards

A softswitch solution emphasizes open standards as opposed to the legacy Class-4 switch that historically offered a proprietary and closed environment (see *Figure 5*). A carrier was a "Nortel shop" or a "Lucent shop." No components (hardware or software) from one vendor would be compatible with products from another vendor. This usually results in less competitive pricing for those components. Softswitch open standards are aimed at freeing service providers from vendor dependence and the long and expensive service-development cycles of legacy switch manufacturers.

The absence of a robust feature list identical to those found on a DMS-250, for example, is an objection among service providers. Softswitch offers the advantage of allowing a service provider to integrate third-party applications or even write their own. While softswitch may have lacked a feature set directly comparable with Class 4 in the past, recent developments in the industry indicate that an ambiguous feature set might be an advantage rather than a performance deficiency. This is potentially the greatest advantage to a service provider presented by softswitch technology.

Softswitch Is a Disruptive Technology

Softswitch enables a lowering of the barriers to entry for service providers in the converging long-distance market. Compared to the Class-4 switch, softswitch, in the form of a Class-4 replacement, is a disruptive technology. Christensen's *The Innovator's Dilemma* sets parameters to compare and contrast the Class 4 and softswitch in the converging market.

Softswitch has a number of technical deficiencies that hinders its growth in market-share in the converging market. They are scalability, reliability, QoS, and well-defined features and applications. Softswitch lacks the scalability of the Class 4 in that it does not yet match the density or call processing power of a Class-4 switch. Softswitch is not yet established as consistently delivering "five nines" of reliability.

Softswitch must overcome QoS issues to win trust in the converging market. Engineering improvements in the gateways and IP networks can accomplish this. Some media gateways can deliver voice quality that is indistinguishable from the PSTN.[12] The Class-4 switch offers service providers and subscribers many standard applications and features. The softswitch industry has yet to succinctly define applications and features, but its open architecture potentially allows it to offer *more* features than Class 4. Softswitch vendors are striving to correct the aforementioned deficiencies, and when they do, softswitch will replace Class 4 in the converging market.

Just as softswitch has quantifiable deficiencies, it also has a number of quantifiable advantages over Class-4 technology. Those advantages are that it is cheaper and smaller and offers distributed architecture and open standards.

FIGURE 4

Softswitch Replaces Class 4 in the Converging Market

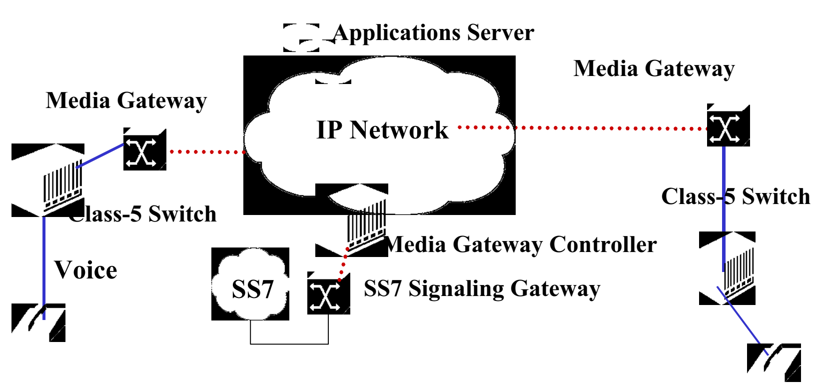

Distributed media gateways installed in or near COs would negate the need for large, centralized Class-4 switches in a converging market.

FIGURE 5

Three-Layer Softswitch Architecture Compared with Legacy Circuit Switching

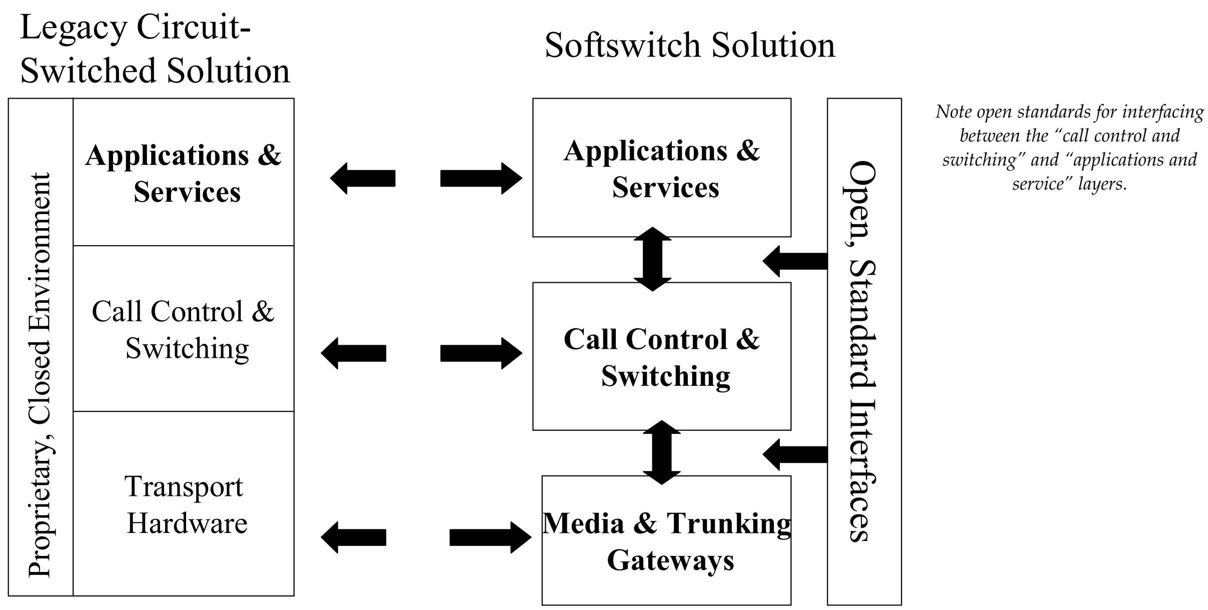

It is cheaper than a Class 4 in both purchase and OAM&P costs. A discrete form factor smaller than a Class-4 switch coupled with distributed architecture gives softswitch the advantage of rapid roll out in markets that cannot be quickly or economically serviced by Class 4. Softswitch offers open standards as opposed to Class 4's proprietary standards. Open standards enable softswitch to offer expanded high-margin applications and services.

It is unlikely that incremental changes to the Class 4 will satisfy the converging market. A more likely scenario is that the incremental improvement of softswitch technology will over come the performance deficiencies of scalability, reliability, QoS, and features and applications while maximizing the advantages of its technology. As this occurs, softswitch will displace Class 4, making softswitch a disruptive technology.

Forecast

Using the disruptive innovation thesis, it is important to compare an existing technology set, the Class-4 switch, with a new set, the Class-4 replacement softswitch, in the converging market. Softswitch is deficient compared to Class 4 in reliability, scalability, QoS, and features and applications. Softswitch has some new capabilities also measurable (distributed architecture, lower cost, smaller form factor, and open standards) that appeal to the converging market in preference to the Class 4.

Either the Class 4 evolves to technology performance requirements of the converging market or softswitch, over time, is incrementally improved until its original performance deficiencies are erased while it still retains its advantages. In this way, the softswitch disrupts Class 4.

Conclusion

Softswitch, despite having technical deficiencies relative to the Class-4 switch, has enabled a number of long-distance service providers to lower the barriers to entry into the converging long-distance market. The advantages of the softswitch—being that it is cheaper and smaller and offers distributed architecture and open standards—made its initial entry in the converging market possible. In the converging market, softswitch vendors are correcting the technical deficiencies suffered by the softswitch regarding scalability, reliability, QoS, and features and applications. By virtue of competing in the converging market, being initially technically inferior to the mainstream Class-4 switch, and having certain advantages over the Class 4, softswitch is a disruptive technology that will replace the Class-4 switch.

Notes

1. ESS No. 4 built for AT&T. The first ESS (No. 1) went into service in 1965.
2. In 1976, there were no personal computers (PCs) and no public access to the Internet. The telecommunications business climate was pre-divestiture and pre-Telecommunications Act of 1996.
3. Robert Chapuis and Amos Joel, "In the United States, AT&T's Digital Switch Entry No. 4 ESS, First Generation Time Division Digital Switch" in *Electronics, Computers and Telephone Systems* (North Holland Publishing, 1990), 337.
4. Christensen, Clayton M., *The Innovator's Dilemma* (Harper Business, 2000, reprinted by arrangement with Harvard Business School Press), xviii.
5. "Introducing the International Softswitch Consortium," a presentation by the International Softswitch Consortium, http//www.softswitch.org.
6. John C. Eline and Martin Pyykkonen, *New Softswitch Technology: The Next Evolution for Today's Public Telephone Network* (C.E. Unterberg, Towbin, June 25, 2001), 6.

7. "Product Service Information – DMS300/250 System Advantage," Nortel Networks (2001), http://www.nortel.com.
John C. Eline and Martin Pyykkonen, *New Softswitch Technology: The Next Evolution for Today's Public Telephone Network* (C.E. Unterberg, Towbin, June 25, 2001), 28.
8. "Accelerating the Deployment of Voice over IP (VoIP) and Voice over ATM (VoATM)," white paper from Telica as posted by International Engineering Consortium, http://www.iec.org/.
9. John C. Eline and Martin Pyykkonen, *New Softswitch Technology: The Next Evolution for Today's Public Telephone Network* (C.E. Unterberg, Towbin, June 25, 2001), 29.
10. Krapf, Eric, "VoIP Services Emerging," *Business Communications Review* (January, 2000), 17; ttp://www.bcr.com/bcrmag/2000/01/p17.asp.
11. Craig, Andrew, "Qualms of Quality Dog Growth of IP Telephony," *Network News* (November 11, 1999), 3.

Media Gateway Control Protocol: Providing Structure in Today's IP–Based Telephone Systems

Randy Salminen

Senior Systems Engineer

Integral Access, Inc.

Executive Summary

The Media gateway control protocol (MGCP) provides the call control necessary to allow wide-scale Internet protocol (IP) telephony over the Internet with interoperability between legacy telephone networks. Efficient and suited for growth and interoperability with legacy voice connections, MGCP is the number-one option for providing signaling over IP networks today.

1. Introduction

Voice over IP (VoIP) is here. Only a few short years ago, we learned about this packet-based tool called the Internet. This breakthrough technology initially enabled us to communicate data between remote locations; later, it allowed us to explore sites for knowledge, shopping, and entertainment. As technology evolved, graphics became more intense and lifelike. High-resolution graphics then led the way for video clips. Now we can even watch and participate in live voice and video connections across this worldwide packet-switched network.

1.1. VoIP Development

Unlike its circuit-based cousin (legacy telephone network), the Internet is packet-based. It was initially designed for worldwide data exchange but has outgrown the concepts and imaginations of its creators by leaps and bounds. The Internet is here to stay (and grow) for a very long time. In addition, VoIP will continue to evolve and grow as an important part of the Internet and other packet-switched networks. However, in order for both voice and video traffic to travel over a data-oriented network and arrive without any noticeable delay in sight or sound, improvements were necessary.

1.2. The Future of VoIP

Since IP and the Internet are ubiquitous, and because the existing public switched telephone network (PSTN) is extremely inefficient in its use of bandwidth, improvements are continually being made. Equipment vendors are quickly implementing protocols such as multiprotocol label switching (MPLS) to provide quality of service (QoS) and turn the Internet into a truly real-time, user-friendly atmosphere. As a result, VoIP is one of the fastest growing industries today (see *Figure 1*).

1.3. Why We Need MGCP

In order for the Internet to perform the functions of a full-fledged voice network (and beyond), it needs a structure that allows a user to make a call to any phone anywhere—whether the call is from an IP phone to a legacy phone, or vice versa. Of course, the resiliency of cellular phones, as well as the quality and features of landline-based legacy phones, is needed. MGCP provides the structure required in an IP–based telephone system.

1.4. What Is MGCP?

MGCP, an Internet Engineering Task Force (IETF) draft, has been chosen as the initial protocol to provide the call control necessary to allow wide-scale telephony over the Internet. MGCP also provides the interoperability between the Internet and legacy telephone networks. MGCP is American Standard Code for Information Interchange (ASCII)–based and resides in a user datagram protocol (UDP) packet, relying on a positive acknowledgment mechanism to deal with packet loss. As its basis, MGCP's stateless protocol was designed using a combination of Cisco Systems and Bellcore's simple gateway control protocol (SGCP) and the Level 3 Technical Advisory Committee's Internet protocol device control (IPDC) specification. Combining these two protocols facilitated a single specification, and, in turn, enabled faster deployment. To ensure security, MGCP uses the IP security (IPSec) protocol and has the ability to use encryption on the media streaming at the terminal equipment. MGCP works with several other protocols to provide the various functions that will be needed in the Internet for VoIP.

FIGURE 1

Internet Uses Today

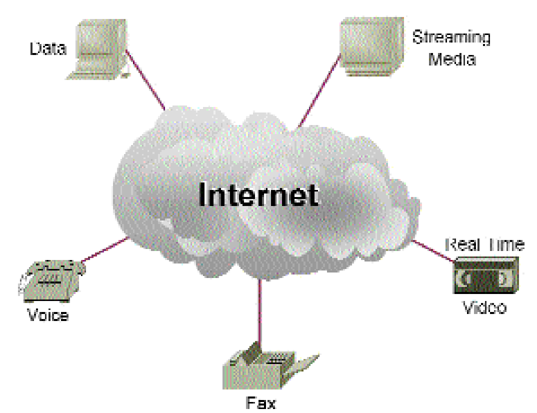

2. Software Architecture

MGCP is used to communicate between a call agent and a gateway. Together, these help to make up the softswitch architecture (see *Figure 2*). The call agents (also referred to as media gateway controllers) are the brains of the operation, and the gateways are the muscle. You need both of them to complete a call over a public network using MGCP. The call agent sends commands to the gateway and receives an acknowledgment for each one. An advantage to call agent/gateway architecture is that new services can easily be introduced without having to swap or upgrade gateways. This model appears as if it is a single gateway, but it is actually distributed. This is commonly referred to as the decomposed gateway architecture. Putting the call-control functions in the call agent and the media streaming functions in the gateway allows for both greater scalability as well as more simple, inexpensive, and reliable local access equipment.

2.1. The Call Agent

The call agents, also called media gateway controllers or gatekeepers, are scattered about the Internet and act as telephony routers directing calls from place to place. The call agent is required for signaling within a call but is not involved in passing the voice traffic. Each call agent, responsible for a number of gateways, accepts event messages from them and tells them how to process calls via a relatively small set of MGCP messages and procedures. The call agents do not communicate with each other via MGCP. Call agents communicate via session initiation protocol (SIP) and H.323 (see Section 5 of this paper). Ideally, network providers will set up several gateways within their network to provide access for their users or customers. One or more call agents, or access to a call agent from another provider, should be set up. The call agents need to have the ability to route calls to and from the PSTN as well as the ability to route calls to and from customers of other network providers across other call agents and gateways. Call agents may also route calls to application or media servers.

2.2. The Gateways

A gateway is a device that is designed to join unlike elements (see *Figure 3*). A media gateway joins unlike media. In the case of MGCP, signals are joining or converting media between circuits and packets, and voice and packets. In MGCP, there are a number of different types of gateways performing different tasks. Each of these gateway types has a distinguishing name and allows voice to be encapsulated in packets. The gateways are primarily concerned with audio signal conversion and media streaming, while the call agents deal mainly with signaling functions. The gateways may be positioned at the customer premises where they can be directly connected to plain old telephone service (POTS) telephones, a private branch exchange (PBX), or a data network. They can also be found at a point of presence (POP) or co-location (colo) where they can service larger numbers of users as well as provide other high-level tasks. Gateways can be associated with multiple call agents to provide redundancy.

2.2.1. The Residential Gateway

The residential gateway may be found in the home or office. Its purpose is to connect the IP telephony subscriber to the packet network and accept events from said subscriber, passing them to the call agent for processing. This gateway also accepts commands from the call agent to set up, control, and tear down calls to and from the subscriber. Residential gateways may be digital subscriber line (xDSL), cable modems, and even broadband wireless devices.

2.2.2. The Trunking Gateway

The trunking gateway, which sits between the PSTN and the Internet, handles the conversion between a legacy telephone

FIGURE 2

The Basic Softswitch Architecture

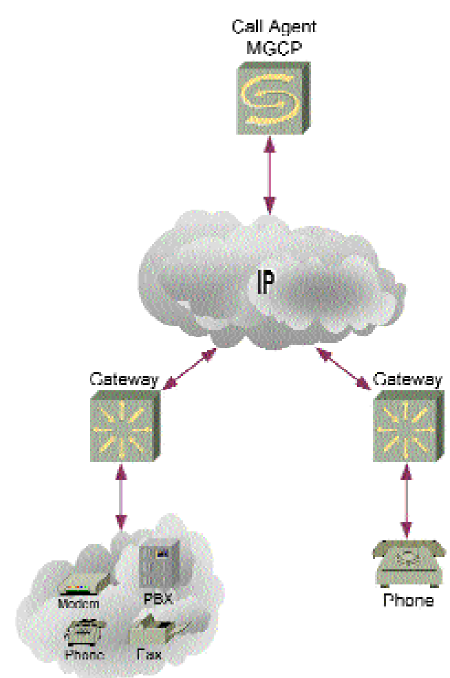

FIGURE 3

The Gateway Model

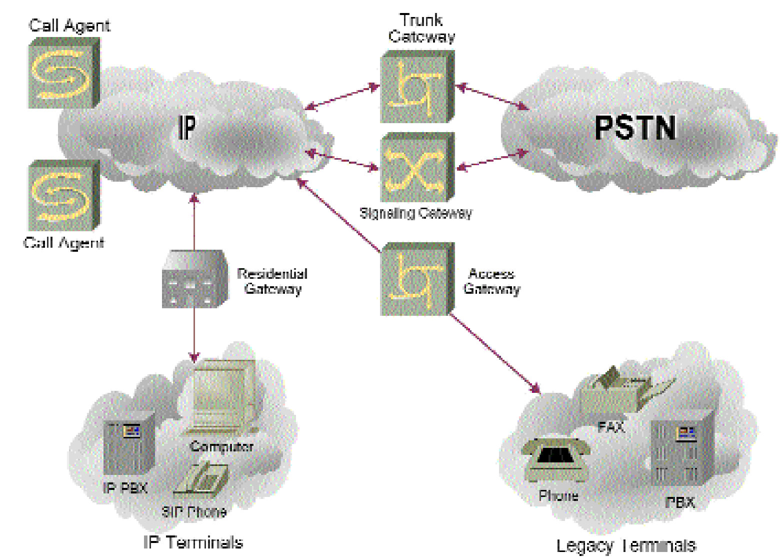

network and an IP network. The trunking gateways are large capacity devices capable of handling tens of thousands of digital signal (DS)–0 circuits, which they transform from time division multiplexed (TDM) voice channels into real-time transport protocol (RTP) or MPLS packets.

2.2.3. The Signaling Gateway

The signaling gateway is designed to handle the signaling conversion and forwarding between the PSTN and a packet-switched network. It has large capacity like the trunk gateway and supports a variety of signaling protocols, including signaling system 7 (SS7) and C7.

2.2.4. The Media Access Gateway

The media access gateway is used to aggregate end user traffic and call-control signaling. Operating under call-agent control, the media access gateway originates and terminates packet-based voice services, class features, and emerging multimedia communication services. While it is technically possible for call agents to directly control thousands of customer-premises devices directly, it is logically impractical. Aggregating MGCP sessions on a media access gateway vastly simplifies overall network management and offers the following advantages:

- Prevents malfunctioning, malicious, or over-enthusiastic customer-premises equipment (CPE) from overwhelming the call agent to have a negative impact on network performance and potential denial of service to others
- Provides a throttling mechanism for varied levels of service to others
- Provides a signaling aggregation to the call agent
- Helps prevent toll fraud, a significant issue for a service provider that wants to make money

2.3. The Connection Model

MGCP uses a connection model based on endpoints, calls, and connections. The connection model is a way to reference points and functions that are used to make a connection.

2.3.1. Endpoints

The term *endpoint* is used in MGCP to describe a physical or virtual port on a gateway or call agent (see *Figure 4*). Some

FIGURE 4

Endpoints

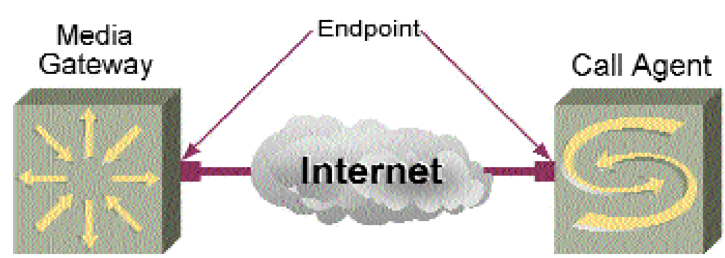

examples of endpoints are the termination point for a DS–0, an Ethernet, an audio server, or a POTS line. Media enters or exits through an endpoint. MGCP addresses these endpoints via endpoint identifiers.

2.3.2. Connections

A connection is an association of endpoints for the purpose of passing traffic between them (see *Figure 5*). Connections are typically formed dynamically, as in the case of a phone call. They are point-to-point, or point-to-multipoint. A multipoint connection is an association of multiple endpoints allowing for conference calls. Connections can be made between separate gateways or between different endpoints on the same gateway. They also have mode settings such as send-only, send-receive, and receive-only.

2.3.3. Calls

A call is a logical association of connections between two or more endpoints (see *Figure 6*). A three-way call contains two connections between three endpoints. Calls are managed by the call agent using domain names to identify each one individually. The call agent gives the call a name when it is created to allow greater versatility. The use of addresses as identifiers may restrict mobility.

2.3.4. Events

Events are occurrences in endpoints (e.g., a telephone goes off-hook). A call agent may ask to be notified of certain events and may also incur events such as telling a gateway to supply dial tone to an endpoint. Events are identified by names and are organized into differing name groups called packages. A package contains events for a specific type of

FIGURE 5

Connections

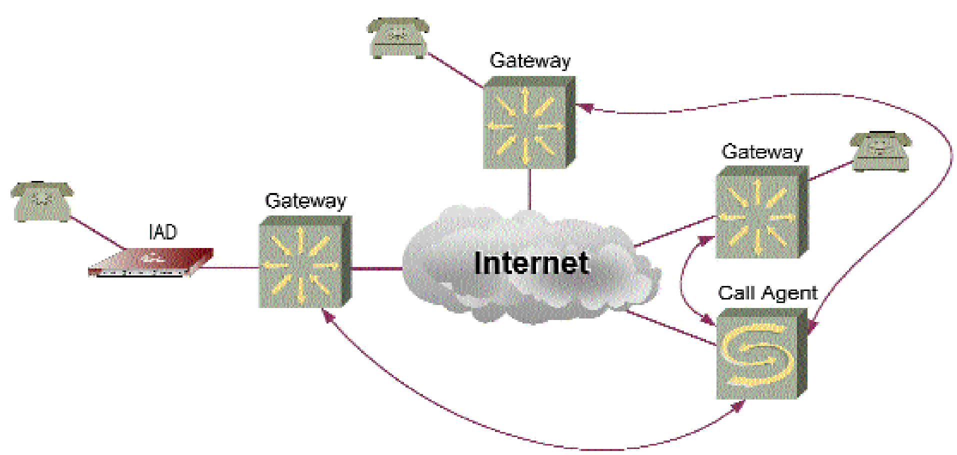

FIGURE 6

Calls

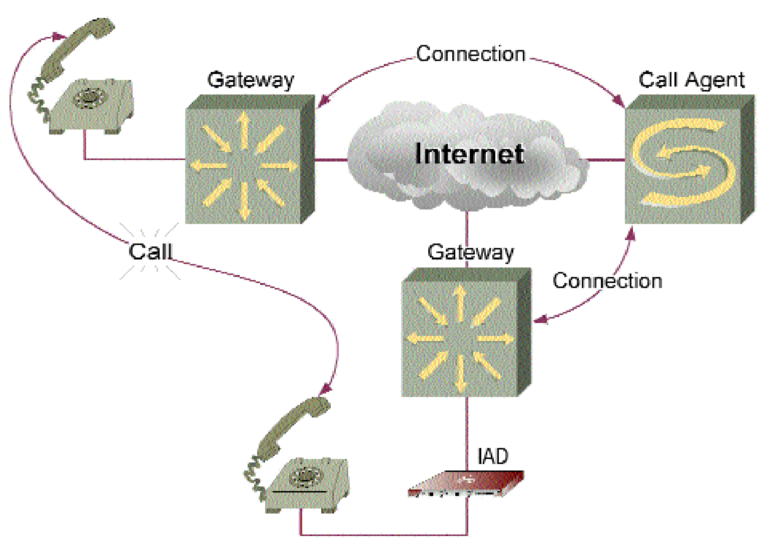

endpoint, such as an analog line. The names, made up of strings of letters, hyphens, and digits, have two parts: event and package.

3. MGCP Commands

Commands are issued from the call agent to the gateways for setting up and tearing down connections. The gateway sends back responses in return. Each command and response must be acknowledged by the receiver to ensure that none will be missed. Commands are also issued to keep the call agent informed of events at the gateways. There are eight commands defined in MGCP:

1. NotificationRequest
2. Notify
3. CreateConnection
4. ModifyConnection
5. DeleteConnection
6. AuditEndpoint
7. AuditConnection
8. RestartInProgress

3.1. NotificationRequest

The NotificationRequest is sent by the call agent to request notification about a particular event.

Example: The call agent wants to be notified when an off-hook occurs on an endpoint connected to a phone. Notification is necessary because an off-hook is an indication that a call will be taking place and that the call agent needs to take action. The call agent will tell the gateway to issue a dial tone and collect the dialed digits.

3.2. Notify

The Notify command is issued by the gateway in response to a NotificationRequest. The gateway includes a list of the events that it observed in the Notify. It also includes the identifier sent in the NotificationRequest so that the call agent understands the origination.

Example: A Notify could include the dialed digits that came from the phone attached to that endpoint.

3.3. CreateConnection

A CreateConnection is used to build connections on and between endpoints for the purpose of making calls. A CreateConnection is sent to each gateway involved, at which time an IP address and UDP port is associated with the endpoint. A CallId is specified to tie the connections to a call. More than one connection may share a CallId to accommodate situations such as when there is call waiting or conference calling. The CreateConnection includes both the parameters for the connection and the mode of the connection.

Some parameters include the following:

- Voice encoding
- Silence suppression

Some modes include the following:

- Send
- Send/receive
- Receive
- Conference
- Inactive
- Data
- Loopback
- Network loopback
- Continuity test
- Network continuity test

3.4. ModifyConnection

The call agent uses the ModifyConnection to change parameters in existing connections. The parameters used in the ModifyConnection are identical to the parameters in the CreateConnection. The ModifyConnection can also be used for activating and deactivating a connection and for providing information about the other end of a connection.

3.5. DeleteConnection

The DeleteConnection can be sent by the call agent or the gateway to delete an existing connection. The call agent may send the DeleteConnection because the other end of the connection has hung up, or it can be sent by the gateway if it can no longer sustain the connection for any reason. The gateway responds to a DeleteConnection by sending a list of parameters back indicating the final status of the connection.

The parameters include the following:

- The number of packets and octets sent and received
- The number of packets lost
- Inter-arrival jitter
- Average transmission delay

3.6. AuditEndpoint

An AuditEndpoint is a request from the call agent for details about the status of an endpoint or a number of endpoints. The gateway sends a response with audited information including the following:

- Requested events
- Dialing plan
- Connection identifiers

3.7. AuditConnection

The AuditConnection is sent by the call agent requesting information about a connection. The types of information that can be requested include the following:

- Call ID
- Local and remote connection descriptors

Local connection parameters and the mode of the connection

The gateway sends back a response with the information requested.

3.8. RestartInProgress

The RestartInProgress is used by the gateway to indicate the type and to tell the call agent that an endpoint or a group of endpoints has been put in service or taken out of service. The choices are as follows:

- *Graceful:* Indicates that the endpoint will be taken out of service after a slight delay
- *Forced:* Indicates that the endpoint is being taken out of service immediately

- *Restart:* Indicates that the endpoint will be placed in service, possibly after a specified delay

4. A Typical Call Flow

To understand MGCP in practice, we have to walk through a call.

1. A user decides to place a call.
2. The sending user picks up the phone and hears dial tone.
3. The sending user dials the number and hears the ringing.
4. The receiving phone rings to indicate an incoming call.
5. The receiving user picks up the phone and immediately hears the sender.
6. Sender and receiver converse.
7. When they are done, one of them will hang up the phone.
8. The other party then hangs up the phone.

4.1. What Happens behind the Scene

Now let us see what is happening behind the scene. We will assume the user placing the call is using a phone connected to the Internet and that the receiver is connected to a phone on the PSTN. This means the calling party is behind an access gateway, a trunking gateway, and a signaling gateway connecting the Internet and the PSTN.

1. The first message is a RestartInProgress, which is sent by the access gateway, indicating the availability of its endpoints.
2. Before the phone goes off-hook, the call agent had sent a NotificationRequest to the access gateway, telling it to monitor for off-hook conditions.
3. The gateway acknowledges the command.
4. The calling party picks up the phone.
5. The off-hook condition causes the gateway to send a Notify to the call agent.
6. The call agent sends back an acknowledgment.
7. A NotificationRequest is sent by the call agent, telling the gteway to put out a dial tone and collect digits.
8. The gateway responds with an acknowledgment.
9. Dial tone is put out on the line connected to the calling party.
10. The calling party dials the phone.
11. The gateway sends a Notify to the call agent with the dialed digits.
12. The call agent sends the gateway an acknowledgment.
13. The call agent then sends a NotificationRequest to the gateway telling it to stop collecting digits and to monitor for an on-hook condition.
14. The gateway acknowledges the call agent (see *Figure 7*).
15. The call agent sends a CreateConnection to the gateway to create the connection.
16. The gateway sends an acknowledgment that contains all of the necessary information about the endpoint.
17. The call agent determines where to route the call, then sends a CreateConnection to the receiving gateway.
18. The gateway sends an acknowledgment along with its relative information.
19. The call agent then sends an SS7 ISUP initial address message to the signaling gateway that will allow the signaling to be hooked into, and understood by the PSTN.
20. The signaling gateway forwards the message to the proper local-exchange carrier (LEC).
21. The PSTN returns a set-up complete message to the Signaling Gateway.
22. The signaling gateway forwards the message to the call agent.

FIGURE 7

A Call Set-Up—Phase One

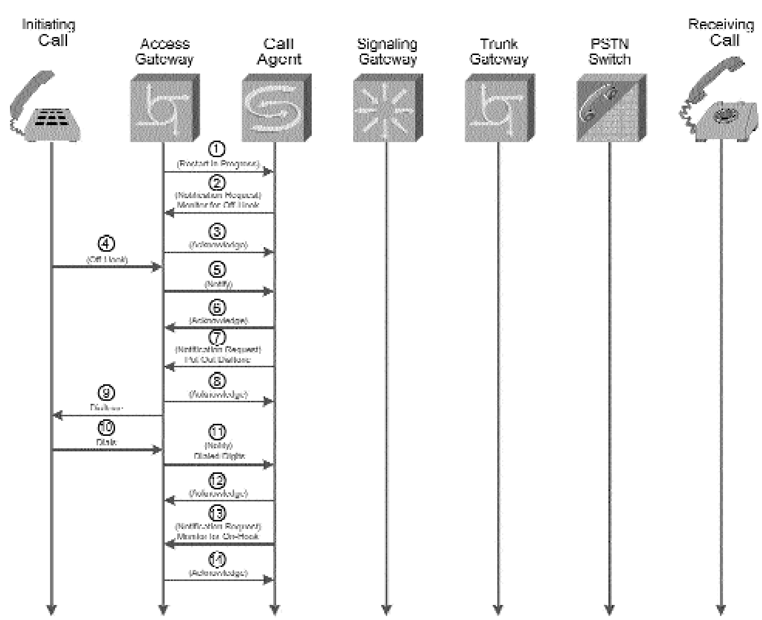

23. The call agent sends a NotificationRequest command to the access gateway to tell it to start ringing.
24. The gateway acknowledges the command and puts ringing on the line.
25. The called party picks up the phone.
26. The PSTN sends an ISUP answer message to the call agent.
27. The call agent sends a NotificationRequest message to the access gateway telling it to take off the ringing.
28. The gateway sends an acknowledgment to stop ringing.
29. The call agent sends a ModifyConnection command to the gateway to put it into a send/receive mode.
30. The gateway sends an acknowledgment and goes into send/receive, and the call is established. The parties can now converse (see *Figure 8*).
31. When they are done, one of the parties will hang up first. Assume it is the called party.
32. A release message gets sent to the call agent from the LEC.
33. The call agent sends a DeleteConnection to the gateways.
34. The gateways will respond with an acknowledgment and include the final status of the connection.
35. The calling party hangs up the phone.
36. The access gateway sends a Notify to the call agent.
37. The call agent sends a NotificationRequest back to the gateway, telling it to monitor for off-hook.
38. The access gateway sends back an acknowledgment that the call is finished (see *Figure 9*).

This is how MGCP works.

5. Other Signaling Protocols

In addition to MGCP, there are a number of other protocols that are possible alternatives or complements to MGCP. Two of the most prominent are H.323 and SIP.

5.1. H.323

H.323 is an umbrella protocol consisting of a number of different sub-protocols that are used to support audio, video, and data applications. There are three versions of H.323:

- The ITU specified version 1 in 1996.
- Version 2 in 1998.
- Version 3, the mature signaling protocol

Version 3 shall be used for this discussion.

The H.323 suite was designed to do the following:

- Provide video, audio, and data across various network types
- Do point-to-point and multipoint conferences
- Handle call control, multimedia environments, and prioritization
- Function across a local-area network (LAN) or wide-area network (WAN)

Its model is similar to MGCP in that it uses gateways to do the media conversion and gatekeepers (the H.323 term for

FIGURE 8

A Call Set-Up—Phase Two

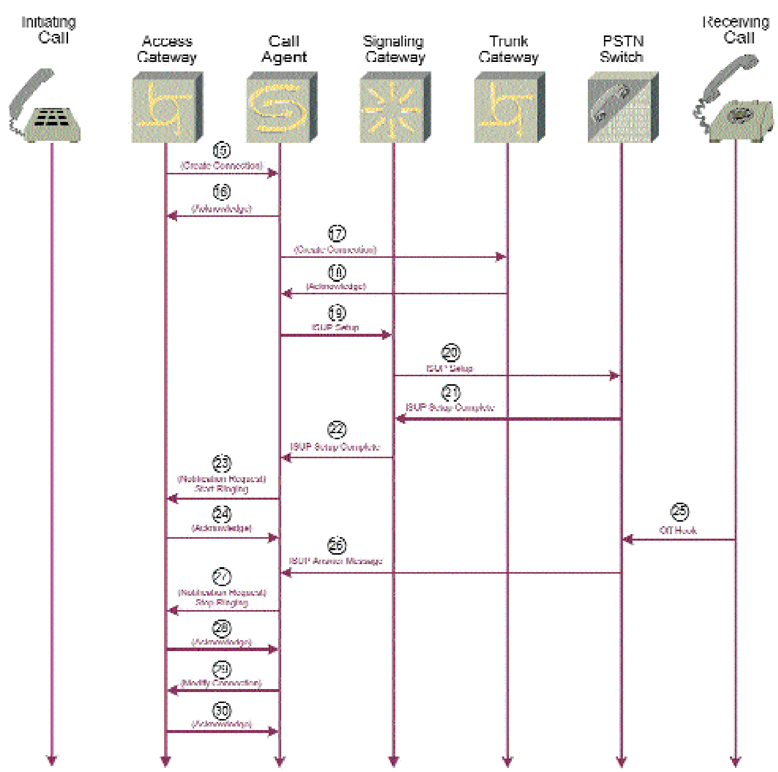

FIGURE 9

A Call Set-Up—Phase Three

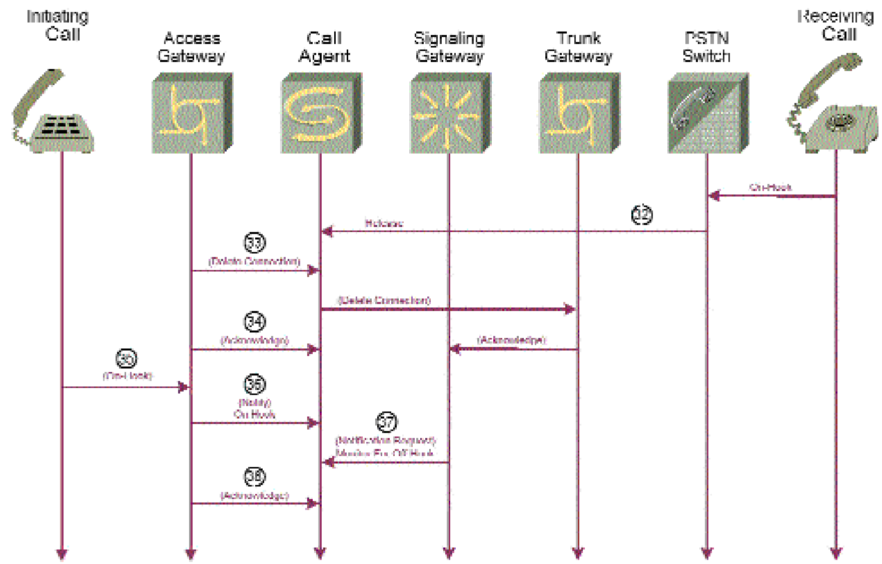

call agents) to do the signaling. The main difference, however, is that unlike MGCP, H.323 was designed to work primarily with intelligent endpoints (such as IP–based telephones), whereas MGCP was intended to function with integrated (intelligent) access devices (IADs) and simple, dumb, endpoints (such as a normal telephone).

H.323 is a complex protocol with a heavy code base. This makes it a very intelligent protocol, but it also causes it to be more expensive to build gateways, and it makes it more difficult to build a simple, reliable gateway. H.323 has a concentrated peer-to-peer architecture that does not allow easy scaling—a key concern to large network providers that need millions of connections. Other shortfalls include the following:

- H.323 is based on transmission control protocol (TCP)/IP, which can cause too much delay to make it practical.
- It will connect over LANs that have no QoS control—an unacceptable solution when you are connecting to the PSTN (which does not tolerate delays).
- Neither H.323 nor SIP was designed to interface with the PSTN—a drawback to using either as a full-scale solution.

5.2. SIP

SIP is a signaling protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. SIP is a text-based protocol, similar to hypertext transfer protocol (HTTP) and simple mail transfer protocol (SMTP), for initiating interactive communication sessions between users. Such sessions include voice, video, chat, interactive games, and virtual reality. SIP is basically the next generation of H.323. It is less complex and better designed than H.323 and is well suited for connection set-up and media mapping. It is designed similarly to H.323, as it communicates with intelligent endpoints. SIP was brought about to dethrone H.323, but since H.323 is pretty well entrenched, most likely the two will coexist for some time.

5.3. H.248/MEGACO

H.248 is the name given to this protocol by the International Telecommunication Union (ITU), and MEGACO is how it is recognized today by the IETF. This is the first time that these two groups have worked on a standard together and is indicative of the future meshing of telecommunications and the Internet. H.248 has a lot of similarities to MGCP and is typically referred to as the protocol that will eventually supersede it. H.248 uses the gateway/call-agent architecture and also handles the call set-up much like MGCP; however, the naming of the parts and functions differs. H.248 also adds enough complexity to allow it to support video traffic (see *Figure 10*). Though there has been successful testing of H.248, it is still an infant in the world of signaling protocols.

6. Surrounding Protocols

As stated earlier, MGCP is a protocol that works together with others to perform the feat of delivering voice over packets. The protocols that surround it are many and varied. We will touch on several of the most important and prevalent surrounding protocols.

6.1. SDP

The session description protocol (SDP) is used to help set up a connection. Defined in RFC 2327, its job is to hand the endpoints the information needed to be compatible with the network (i.e., the type of call being made and its opposite endpoint). The call agent that is associated with the gateway on which that endpoint is located uses SDP. SDP commands will be sent from the call agent at the time of call set-up or modification. Some examples of descriptors are IP addresses, payload types, and RTP port numbers.

FIGURE 10

The Complete Solution

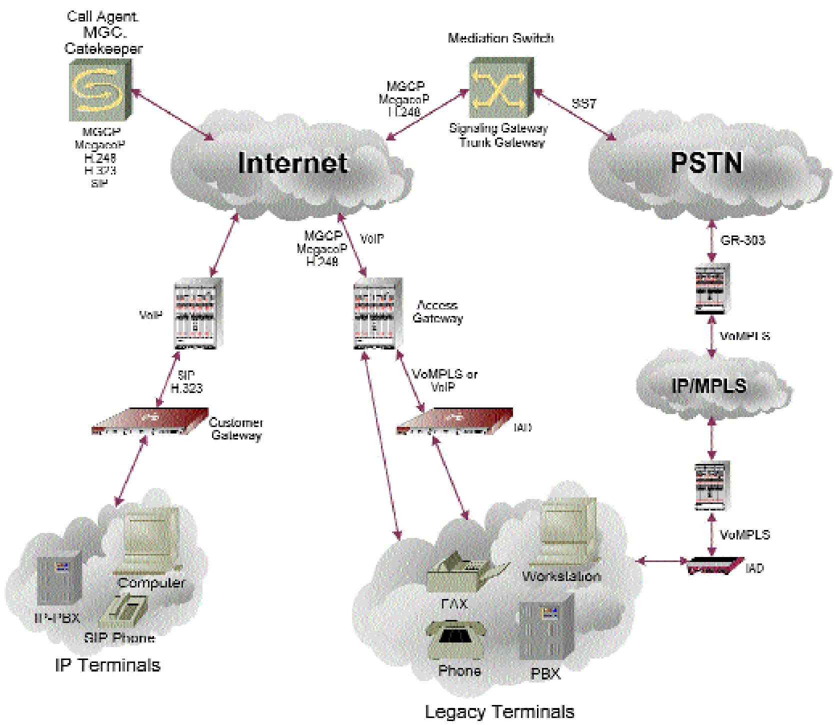

6.2. RTP

RTP is used for transporting real-time audio or video streams over packet-switched networks. MGCP is only used to set up, modify, and tear down the calls; it needs a protocol like RTP to handle the actual voice traffic. RTP does not reserve bandwidth; it can only send information to the receiver. The receiver has to be able to act on the information through things such as jitter buffers. Some of the information provided by RTP includes the following:

- Time stamps, so that the receiver can determine if there is any delay
- Sequencing numbers, so the order of the packets is preserved
- Payload identifiers, used to express what type of encoding is being used for the media

This is important because bandwidth allocation in the Internet may adjust the encoding to conform to the available bandwidth.

6.3. RTCP

The real-time control protocol (RTCP) works with RTP to help ensure QoS. It will send information to users upon request, providing them with information on round trip delay, jitter, and lost packets.

6.4. MPLS

MPLS has emerged as the preferred technology for providing QoS, traffic engineering, and virtual private network (VPN) capabilities on the Internet. MPLS contains forwarding information for IP packets that is separate from the content of the IP header, such that a single forwarding paradigm (label swapping) operates in conjunction with multiple routing paradigms.

The basic operation of MPLS is to establish label-switched paths (LSPs) through the network into which certain types of traffic are directed. MPLS provides the flexibility of being able to form forwarding equivalence classes and also being able to create a forwarding hierarchy via label stacking. All of these techniques facilitate the operation of QoS, traffic engineering, and VPNs.

Once an MPLS enabled network is established, additional routing protocols or extensions of existing routing protocols can be used to activate the MPLS capabilities. For example resource reservation protocol–traffic engineering (RSVP–TE), a control/signaling protocol, can be used to establish a traffic-engineered path through the routed network for high-priority traffic.

Voice is one of the most important types of traffic carried on today's networks. For an IP network to be truly considered "multiservice," it must accommodate voice traffic and its unique transmission requirements. Unlike data traffic, voice traffic needs to travel consistently through the network without being subject to delay or packet re-ordering. A standard IP network, operating on a best-effort basis, is unable to guarantee such preferred treatment. An MPLS enabled network, on the other hand, is able to provide low-latency and guaranteed traffic paths for voice. Using MPLS, voice traffic can be allocated to a forwarding equivalence class that provides the differentiated service appropriate for this traffic type.

7. Summary

MGCP is a required protocol to allow the proper growth of VoIP. Efficient and suited for growth and interoperability with legacy voice connections, MGCP is part of the signaling-protocol evolutionary process.

Prior to MGCP was H.323, but its design limits it to functioning in niche areas. Also, it doesn't perform optimally. Next is SIP—more simple and efficient then H.323. SIP is suited for environments where the end device has the intelligence to interact with the SIP server and where interaction with legacy switches is not required.

The next evolutionary step is MEGACO, or H.248, which is designed to be optimal for both voice and video. Smart MGCP product vendors are preparing for the next step by building their devices to be ready to turn on MEGACO when it is finally deployed.

8. List of References

1. Media Gateway Control Protocol (MGCP) Version 1.0, RFC 2705
2. PacketCable 1.0 Architecture Framework, Technical Report
3. Integral Access, Inc. MGCP Media Gateway Specification
4. *Voice Over IP*, Uyless Black, Prentice Hall Series
5. Rakesh Arora, *Voice Over IP: Protocols and Standards*.
6. Jeff Lawrence, "Integrating IP and SS7 Technologies," *CTI Magazine Compass*.
7. Voice Over IP (VOIP) techguide
8. Sandra Gittlen and Denise Pappalardo, "Voice-over-IP camps wrestle over standards," *NetworkWorldFusion* news.
9. Peter Clegg, "The Road to VOIP," *NetworkWorldFusion*

Developing a Distributed Internet-Based PBL Environment Using Java Beans, CORBA, and JDBC

Mohammed Samaka

Associate Professor of Computer Science and Director of the Computer Science Program, College of Information Technology

United Arab Emirates (UAE) University

Abstract

Problem-based learning (PBL) is a student-centered learning method that provides students with problem-solving, self-directed learning, and collaborative learning skills. PBL has become an increasingly popular method of learning in many professional fields—for instance, in industrial training and in university education that require group management. This research project involves the design, development, and implementation of an instructional software system that runs on the Internet. The system provides a distributed-based multimedia environment that supports the PBL approach. The design is a complete three-tier client/server application for the Web using Java Beans, common object request broker architecture (CORBA), and Java database connectivity (JDBC). It provides synchronous and asynchronous communication services that allow students in one group to communicate remotely either with each other or with the course instructor. This application is being proposed to university and industrial courses that involve projects management, instructor supervision, and group work.

1. Introduction

PBL is an educational format that is centred around discussion and learning. It is a method that encourages independent learning and a deeper understanding of the material rather than superficial coverage (Gallow, 2000). With this approach, students are more involved in their own learning and group work and take responsibility for solving professional problems. The instructor creates a problem situation for the students and then acts as an observer and an advisor while the students work out a solution. A problem used in the PBL approach can be a research project, a case method, or a design project. The small group setting used in PBL encourages an inquisitive and detailed look at all issues, concepts, and principles contained within the problem. The time spent outside of the group setting facilitates the development of skills such as literature retrieval, critical appraisal of available information, and the seeking of opinions of peers and specialists (Thomas, 1995).

The traditional classroom teaching pattern views the teaching process as transmission of knowledge that relates to students' future professional role, while the learning process is an acquisition of that knowledge (Bridge & Hallinger, 1995). This pattern was based on the assumption that learners are able to recognize when it is appropriate to use the acquired knowledge. However, it was noticed that the context in which knowledge is learned has little input on subsequent recall or use in analysing and solving a problem situation. On the other hand, PBL rests on an entirely different set of assumptions. PBL founders assume that the knowledge and the ability to use that knowledge are of equal importance. They also assume that the problems that students are likely to encounter in their future professional practice provide a meaningful learning context for acquiring and applying new knowledge.

PBL is a newcomer to the field of educational administration. It has been used for more than a decade by McMaster University to prepare future physicians and other professionals. It was then incorporated as a curriculum component into a number of institutions in United States and Canada. Nowadays, PBL has become an increasingly popular method of learning in many professional fields as well as in university education (Cameron & Barrows, 2000; Yip, 2000; Pross, 2000).

There is a growing trend for universities to use the World Wide Web (WWW) in direct support of teaching and learning. Various universities around the world are using the WWW as an instructional environment, as both a primary means of information delivery and a supplement to classroom teaching (Harasim, 1999; Hiltz & Wellman, 1997). The use of the Web as a teaching resource can benefit students (Milheim & Harvey, 1998; Sloan, 1997). It is more advantageous when the pedagogical process of PBL is supported by technology to foster teaching and learning. Although learning and teaching courseware is available (Blumenstyk, 1999; Frederickon, 1999) and some universities have been using WebCT as the teaching platform, the available features may not be able to support interactive facilities, various on-line facilities, project planning and controlling, and private discussion facilities for collaborative learning. All are important facilities within PBL.

The project described in this paper concerns with designing, developing, and implementing an instructional application program that supports the PBL approach over the Web. This application is called distributed-based PBL and is abbreviated to DPBL. The DPBL application is an integrated multimedia environment that provides students and instructors with synchronous and asynchronous communication services across heterogeneous computing platforms over the Internet. The DPBL is a complete three-tier client/server application for the Web using Java Beans, CORBA, and JDBC. In this environment, the client beans from the Internet browser invoke operations on CORBA middle-tier server objects via an object request broker (ORB). The server objects provide the application logic, and they store their persistent data in a JDBC–compliant structured query language (SQL) database.

2. DPBL Modeling

During the analysis and modeling processes for the DPBL system, the method of the data-flow model was implemented. This method reflects the end-user understanding of the system and contributes directly to object identification and the identification of operations on the objects. For that reason Oracle-Power Designer-7 was used for composing the data-flow diagrams and the conceptual data module (CDM) for the system. The CDM represents the overall structure and describes the conceptual relationships of different entities and their relations. The power designer was then used to translate the data types specified in the CDM into the physical data types, which the target database supports.

The context diagram in *Figure 1* shows the users that the DPBL system is serving and their relationships over the functions that the system performs. The external entities are students and instructors that are interfaced to the DPBL system through a friendly client interface.

The main DPBL client interface, as shown in *Figure 2*, is a Java frame that is composed of Java Beans. It runs inside a Web browser and provides the friendly interface that allows users to run a PBL facility available in the server via method invocation on the server's object implementations using CORBA infrastructure. The frame shows that two sets of

FIGURE 1

Diagram of the DPBL Application

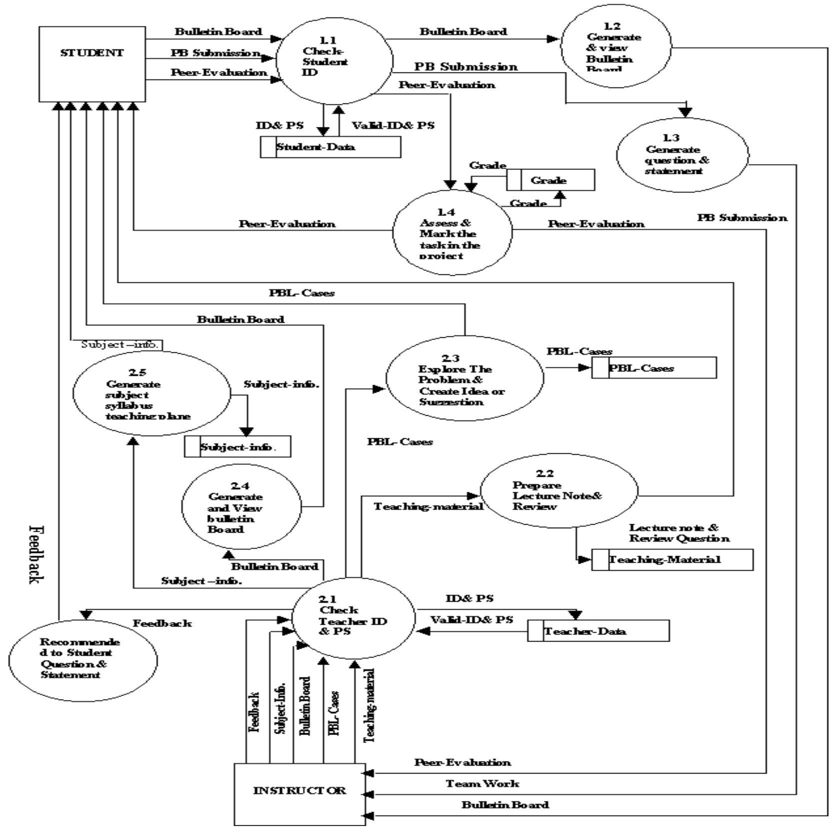

FIGURE 2

The Main Frame of a DPBL Showing the Class Room Facilities

PBL facilities are provided for two Web-based rooms: class rooms and group rooms, respectively.

The DPBL facilities are shown in more detail in *Table 1*. The classroom facilities are provided to every student in the class. Each group of students has a group room. The group-room facilities are accessible only by the members of the group, using a password to log onto the appropriate Web site where the facility is available.

The facilities provided by the DPBL application create the environment for implementing the major PBL components: the realities of the workplace, the content, instructional process, and evaluation. These facilities also sup-

TABLE 1

The PBL Facilities Provided by the DPBL Application

CLASS ROOM	*GROUP ROOM*
Course Material • Course description and learning objectives • Course outline • PBL guidelines	***Problem Material*** • Problem cases • Instructor's notes and guidelines • Relevant references and links
Posting Board • News and posting remarks	***Group Submission Area*** • Project plan • Problem analysis
Instructor Contact Address	***Individual Submission Area*** • Task analysis • Task-completion report
Search Engine	***Interactive Discussion Services*** • Chatting facility • Whiteboard facility
	On-Line Discussion Services • Editing and mailing messages
	Group-Instructor Services • Grading area • A facility for managing and posting course material • A facility for managing and posting student feedback

port the PBL goals in developing students skills in a) problem solving, that is defining, analysing, and solving problems; b) self-directed learning; c) working collaboratively within groups; and d) project management, organizing, and planning.

In the PBL approach, the problem cases that are given to students cover the course learning elements and motivate them to work collaboratively in groups. Using postings on the Web of the information given by the instructor regarding the problem cases, the students will then be responsible for preparing for a meeting session with the instructor. This exercise stimulates the students to focus on the essential elements of the subject and comprehend the subject thoroughly, especially when they revisit the conceptual materials.

Also, each group is responsible for decomposing the problem into its components and assigning a task to each person. Then the group draws up a project-management plan and puts it on the Web. The plan includes task sequence and description, a time schedule, and the allocation of tasks to the group's individuals. This exercise permits the students to learn project-management skills and to play a more active role in their learning. Consequently, their critical-thinking and problem-solving abilities will be further enhanced.

3. Web-Distributed Technology

During the development process, several distributed Web-based learning sites were evaluated to determine which hardware and software might be of use for developing and implementing the DPBL application. For the process of identifying the technologies involved in implementing the DPBL, two main objectives were considered. First, the project has to be restricted to technologies that are available to most potential DPBL students and instructors. Second, students working on different computing platforms across the Internet should be able to use the DPBL. To meet the second objective, a Web-distributed framework should be in use. A Web-distributed framework allows developers to create distributed software that can run across a heterogeneous computing environment through the Internet. It integrates Web computing across different boundaries of different computer machines, operating systems, and programming languages. In commercial sectors, the most prominent frameworks for distributed computing are CORBA, common gateway interface (CGI), distributed component object model (DCOM), and, to a lesser extent, Java RMI. For DPBL development and implementation, CORBA, in particular, was chosen as being the distributed client/server framework that is reasonably suitable. This was based on a previous research project, which involved courseware development that was conducted by the principal investigator of this research work (Samaka & Godfrey, 2000). CORBA is one of the newer technologies in distributed client/server computing. It was introduced and is still being controlled by the Object Management Group (OMG) (Object Management Group, 1996). CORBA now is the leading industry standard for communication between distributed objects. It has already been successfully deployed in a variety of information systems application domains. It allows client and server objects written in different languages and running in different hardware platforms to interoperate.

It also provides operating systems and programming languages independent interfaces. This is because CORBA specifies a language-independent object-oriented architecture and an interface definition language (IDL) for connecting heterogeneous components (Orfali & Harkey, 1998). The IDL compiler generates client/server stubs in the desired target language, which together realise the interconnectivity between an object and remote client.

Furthermore, CORBA provides a suite of high-level services, relieving programmers from having to implement them from scratch, such as, transactional services, security, and a few others.

Finally, it was stated in the literature (Evans & Rogers, 1997; Leippinen & Pulkkinen and Rautianen, 1997) that Java, WWW, and CORBA could be used together in developing highly sophisticated and efficient distributed information systems applications.

4. DPBL Development and Implementation

The DPBL application creates a three-tier client/server environment using Java Beans, CORBA, and JDBC. The client/server environment is chosen because remote communication is required between remote clients over the Internet. In this application, the PBL client beans, which implement the first tier, invoke operations on CORBA middle-tier server objects. The server objects that implement the PBL facilities, in turn, store their persistent data in a JDBC compliant SQL database—Oracle.

The tools used to develop the DPBL system were the Java Development Kit 1.2 (JDK) and Netscape Communicator. The Java ORB used in this development is the Visibroker 4.0 product from the Visigenic. The later product was used for pre-compiling the DPBL IDL interfaces into Java using the idl2java pre-compiler and for providing tools for running the DPBL system over a distributed environment— tools such as OSAgent and the Internet Inter-ORB Protocol (IIOP) with Gatekeeper.

The WebGain Visual Cafe 4 Enterprise edition was used to develop the graphical user interface (GUI) for the DPBL application. Each GUI was implemented using the JFrame component of the Java Swing package. The JFrame is a top-level window that provides the basic attributes and behaviours of a window. Swing components are often referred to as lightweight component, since they provide a uniform appearance across all platforms. On the other hand, the original GUI components, AWI, are called heavyweight, since they are tied directly to the local platform and may cause different appearance in different Java platforms (Deitel & Deitel, 1999).

During the development stage, the front-end frame program was coded in isolation to the server program to reduce the time spent on testing and debugging. The server program was also coded and developed in a similar manner. Once the client frame and the server program worked satisfactorily, CORBA was then implemented between them. The DPBL main IDL interface was coded and pre-compiled into Java using the Idl2java pre-compiler. Also, the function "bind" was included within the frame code to allow client objects to locate objects that implement the DPBL interface, as shown in the following Java statement:

```
try { // Initialize the ORB.
    showStatus("Initializing  the  ORB,  Please
Wait");
    org.omg.CORBA.ORB  orb  =  org.omg.CORBA.ORB.
    init(this, null);
    // Bind to the Course Object
    showStatus("Binding to Course Object");
    course  =  DPBL.CourseHelper.bind(orb,  "My
    Course");
    } catch(org.omg.CORBA.SystemException e){
    showStatus("Applet Exception" + e);
    e.printStackTrace(System.out);
    }
```

To run the DPBL application, you need first to run the server application that requires a Java VM and an ORB (Visibroker 4.0). Once the server application is running alongside the Web server, the client frame can then perform requests on the server objects.

When the client frame is downloaded and bounded to the server, students and instructors can then choose a DPBL facility from the menu. In the DPBL, each facility is normally front ended with a frame, as shown in *Figure 3*.

5. Discussion and Conclusion

By the time this paper was submitted, most the asynchronous facilities provided by the DPBL were implemented and tested. However, further work is needed to complete the other asynchronous facilities (refer to *Table 1*) that provide real-time interactivity between the DPBL participants.

The DPBL application can now be used in university and industrial courses that involve projects, instructor supervision, and a group-work application area. The technology of the application provides facilities to both teachers and students to facilitate learning and teaching. These include posting teaching materials on the Web to support on-line information retrieval, a discussion area for brainstorming, a facility for project planning and monitoring, a private group area for the purpose of discussion, a facility for the submission of development work and learning resources, and a posting area to display good work to motivate students. The students' progress can also be monitored, and their instructor can give feedback to students in regard to their progress.

6. References

1. Blumenstyk, G., "The Marketing Intensifies in Distance Learning," *The Chronicle of Higher Education* (1999), 27–30.
2. Bridge, M., and Hallinger, P., "Problem-Based Learning, Clearinghouse of Educational Management" (1995), http://eric.uoregon.edu/publications/pub_pbl.html.
3. Cameron, T., and Barrows, H., "Maintaining Authentic Problem-Based Learning in a Distributed Learning Environment," Proceedings of the IASTED International Conference of CATE'2000 (Cancun, Mexico: 2000), 76–82.
4. Deitel, H. and Deitel, P., "Java How to Program," (Prentice-Hall, 1999).
5. Evans, E. and Rogers, D., "Using Java Applets and CORBA for Multi-User Distributed Application," *IEEE Internet Computing*, 1(3) (1997), 43–55.
6. Frederickon, S., "Web-Based Instruction: A Comparison of Ten Programs," Proceedings of the 15th Annual Conference in Distance Teaching and Learning (Madison, Wisc.: 1999).
7. Gallow, D., "What Is Problem-Based Learning?," UCI University (2000), http://www,pbl,uci.edu.
8. Harasim, L. "A Framework for Online Learning: The Virtual U," *IEEE Computer* (1999), 44–49.
9. Hiltz, S. R and Wellman, B. "Asynchronous Learning Networks as a Virtual Classroom," *Communications of the ACM*, 9(40) (1997), 44–49.
10. Leippinen, M., Pulkkinen, P., and Rautianen, A., "Java- and CORBA–Based Network Management," *IEEE Computer*, 30(6) (1997), 83–87.
11. Milheim, W. D., and Harvey, D. M., "Design and Development of a World Wide Web Resource Site," *Educational Technology*, 1(38) (1998), 53–56.
12. Object Management Group, "The Common Object Request Broker: Architecture and Specification-Revision 2.0," OMG 96-08-04 (Framingham, Mass.: Object Management Group, 1996).
13. Orfali, R. and Harkey, D., *Client/Server Programming with Java and CORBA*, Second Edition, (Wiley, 1998).

FIGURE 3

An Example of a Frame for Using a DPBL Facility

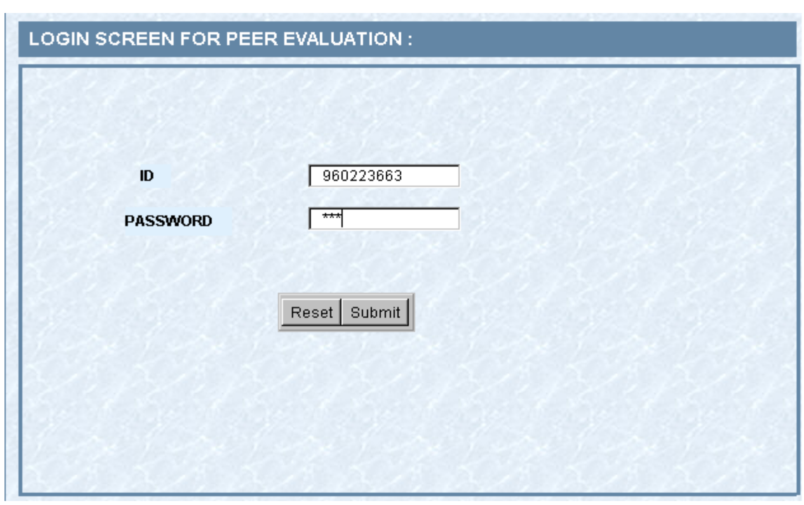

14. Pross, H., "Problem-Based Learning," Queen's University on-line (2000), http://meds.queensu.ca/medicine/pbl/pblhome.htm.
15. Samaka, M. and Godfrey, M. "Implementing Java Applets and CORBA in Courseware Application," Proceedings of the IASTED International Conference of CATE'2000 (Cancun-Mexico: 2000), 321–326.
16. Sloan, A., "Learning with Web: Experience of Using the World Wide Web in a Learning Environment," *Computers and Education*, 4(28) (1997), 207–212.
17. Thomas, J. "A Review of Research on Problem-Based Learning," the Autodesk Foundation (1995), http://www.autodesk.com/foundation.
18. Yip, W., "A Web-Based System for Problem-Based Learning," Proceedings of the IASTED International Conference of CATE'2000 (Cancun-Mexico: 2000), 69–75.

Mesh Fabric: Next-Generation Interconnect for Next-Generation Network Equipment Platforms

Rob Shaddock
Director, Gateway and Wireless Application Unit
Motorola Computer Group

Paul Virgo
Director of Product Marketing, Telecommunications Business Unit
Motorola Computer Group

As the communications market evolves, there is increasing convergence of voice, data, and multimedia applications in telecom networks. Toprovide new applications or services, various types of access devices are being deployed at the edge of the network. The proliferation of these new devices has lead to interoperability and configuration-management issues at the network edge. These issues have in turn escalated the need for a scaleable platform family solution, which can provide powerful processing and storage capabilities for control applications plus a high throughput bearer architecture that supports multiple services on circuit-switched and packet networks.

Telecom equipment manufacturers can realize a significant step forward in network convergence by developing platforms that accommodate differing protocols such as Internet protocol (IP), asynchronous transfer mode (ATM), and frame relay without unnecessary conversion and overhead.

Bearer Protocols

There are numerous protocols within the network that must be addressed by telecom original equipment manufacturers (OEMs) to offer the greatest availability of services for customers. Among the more popular as the following:

- ***TDM (Time Division Multiplexing)*** is the basis for the existing public switched telephone network (PSTN). It uses hierarchies of multiplexed channels with a specific amount of fixed bandwidth. This fixed nature is the biggest drawback for data—rather than voice—transmission, as time slots are allocated based on the amount of users, not on the amount of user traffic being carried.
- ***ATM*** comprises much of the packet-based network backbone. It uses small, fixed-sized packets and supports a variety of different qualities of service routing mechanisms. ATM also supports different types of traffic and optimises network utilisation.
- ***P*** is a simpler protocol than ATM that reduces the cost of network deployment. IP networks are commonly used in local-area networks (LANs) in homes and businesses.

In contrast to traditional networks, packet-based networks use routing and switching to efficiently transport data as demand increases or decreases, rather than using fixed resources based on the number of users. This more efficient use of the network coupled with the much larger amounts of data transport caused by the explosive growth of the Internet are the driving forces behind the need for flexible telecom solutions that address communications between different networks and offer migration to a packet-based world.

Standardization of Platform Architectures

As the industry continues to evolve toward packet-based switching architectures, there has been significant momentum toward standardizing equipment architectures. This standardization will encourage third-party development efforts and enable telecom companies to reach economies of scale for the reuse of technology across various platforms. One of the most significant standards to date that specifically addresses switching fabrics in CompactPCI® systems is PCI Industrial Computer Manufacturers Group® (PICMG) 2.16, which was ratified in September 2001.

The PICMG 2.16 standard for CompactPCI Packet Switching Backplane (CPSB) architectures overlays an embedded Ethernet switching fabric on the backplane in CompactPCI systems. The architecture is based on two concepts: an Ethernet infrastructure is embedded in the CompactPCI backplane and is accessed via the J3 connector, and all subsystems operate as standalone systems on a card, interfacing through a network stack on top of Ethernet. PICMG 2.16 is most useful in the control planes at the edge of the network, where applications are processing intensive, but is also suited to bearer applications that use IP. Typical applications most suited for applications with this type of architecture include voice over IP (VoIP), media gateways, integrated access devices, and embedded service clustering.

A potential use of PICMG 2.16 is to provide an all packet infrastructure with IP as the core packet transport. The architecture provides up to 1 gigabit per second (Gbps) er board of bandwidth for a system capacity approaching 20 Gbps (in large systems). The limitation of this architecture is that the 2.16 standard only supports IP transport, thus limiting its scope largely to metropolitan-area networks (MANs) and LANs. Generally, the core of the network currently uses TDM and ATM for public switching, although there is movement toward gigabit IP transport in the core.

The growth of the converged network has consequently led to several key issues for telecom OEMs:

- How can an OEM efficiently address the high bandwidth, speed, and multiple protocol requirements that occur in networking applications?
- How does a telecom OEM migrate from a traditional circuit-switch network to the all packet networks based on current and emerging technologies?
- How can the OEM leverage the open industry standards regulated by PICMG to ensure that there is the greatest extent of third-party products available for integration into a system-level solution?

We will attempt to answer these critical questions, but first we need to understand the difference in transport topologies.

Star versus Mesh Topologies

There are many topologies for wiring together boards and systems to transport information. They differ in factors such as cost (pins to connect, speed, logic costs, etc.), reliability, and complexity.

Star Topologies (see *Figure 1*) use a point-to-point configuration in which each device uses a dedicated link to send/receive data from a central resource. This resource provides the data distribution for the system. Ethernet networks are a hierarchy of star networks. Each subnetwork is a leg on a star of the next layer in the hierarchy.

Star topologies require redundancy to provide reliability. Reliance on a single central resource can cause a loss of all elements below the failure point. The topology of the PICMG 2.16 specification is a dual star configuration.

Mesh Topologies (see *Figure 2*) are a superset of star topologies and also use point-to-point connections. As interconnects are added to eliminate "dead branches" in a star network, a point is reached where all nodes have connections to all other nodes. At this point, the hierarchy disappears. Each node can be an endpoint, a router, or both.

Meeting the Needs of the New Network

An ideal solution to meet the ever-changing needs of the telecom industry is the creation of an extremely flexible platform that can allow for a variety of protocols and services integrated into a base platform. The infrastructure of the platform must have sufficient bandwidth to meet the current and future needs at the network edge, where the industry will experience significant growth in the foreseeable future.

To meet the high bandwidth and speed requirements of the edge and access portions of the network, telecom OEMs must be able to efficiently process large amounts and different types of traffic.

FIGURE 1

Star and Dual-Star Topologies

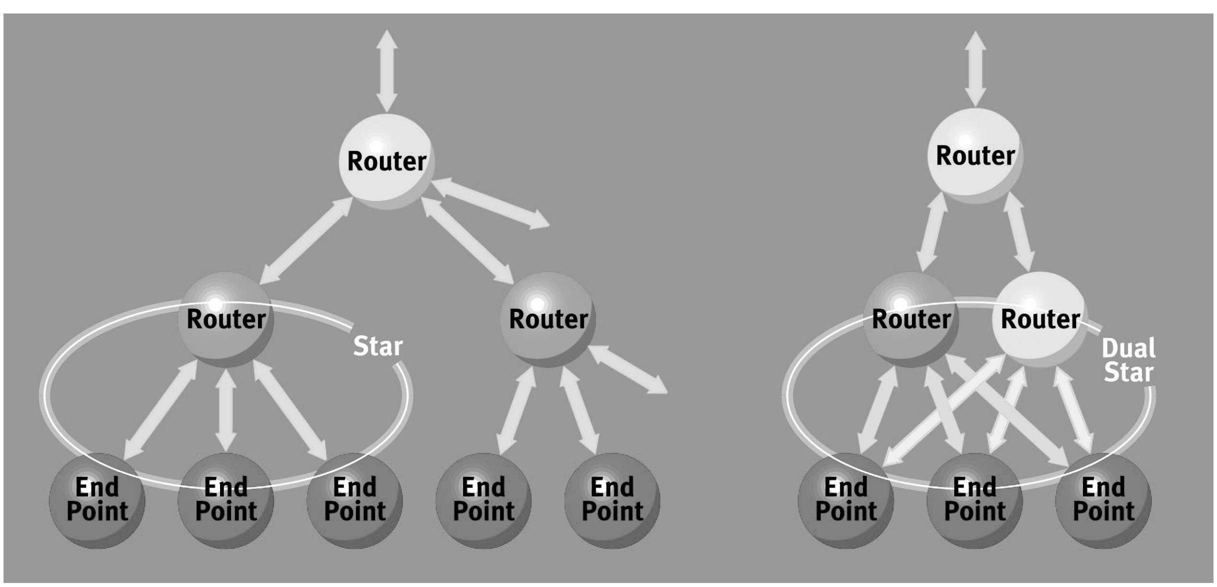

FIGURE 2

Mesh Topology Example

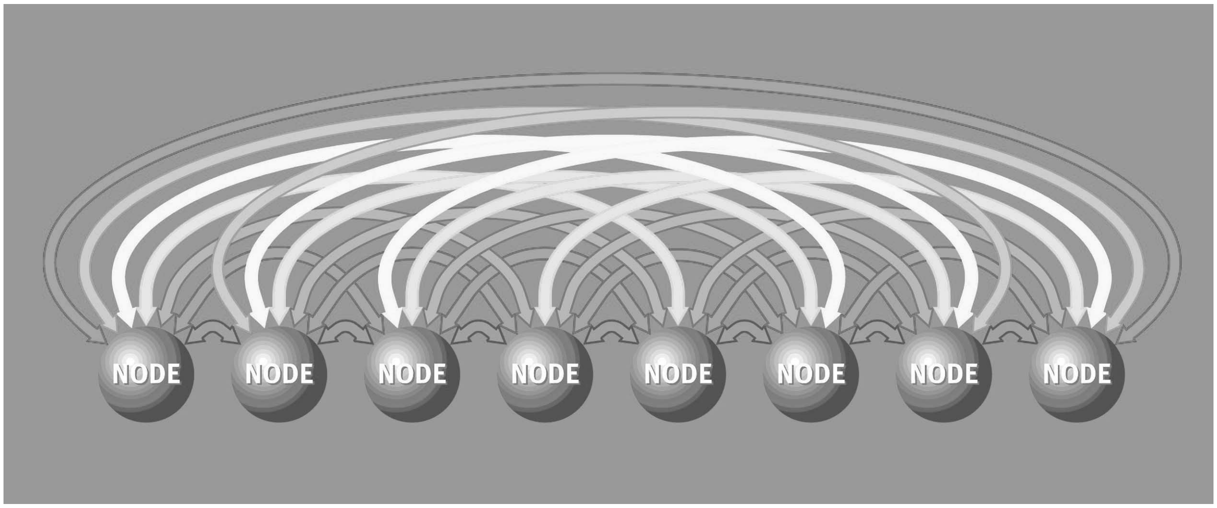

A key strategy to meet these high bandwidth and speed requirements while leveraging the move toward open standards is to add a high-speed mesh data fabric to augment the PICMG 2.16 Ethernet packet switching backplane standard. This fabric could be used for distributed processing, distributed storage, and bearer plane applications.

The high-speed mesh data fabric fills a void not just for bandwidth, but also for protocols and quality of service (QoS) that are not supported by IP. This creates a multicomputing environment that can compete with higher-end symmetric multiprocessing systems.

Combining all of these features in a single data transport platform allows the entire spectrum of edge applications to be addressed (see *Figure 3*). The resulting multi-service platform delivers the best of all worlds—delivering optimum performance in data plane, control plane or integrated data and control applications.

What About TDM?

Clearly, a large portion of the industry has significant investments in circuit-switch networks. To protect these investments, OEMs must seek a solution that combines

FIGURE 3

Adding a High-Speed Mesh Fabric in a Flexible Platform Addresses the Entire Spectrum of Edge Applications

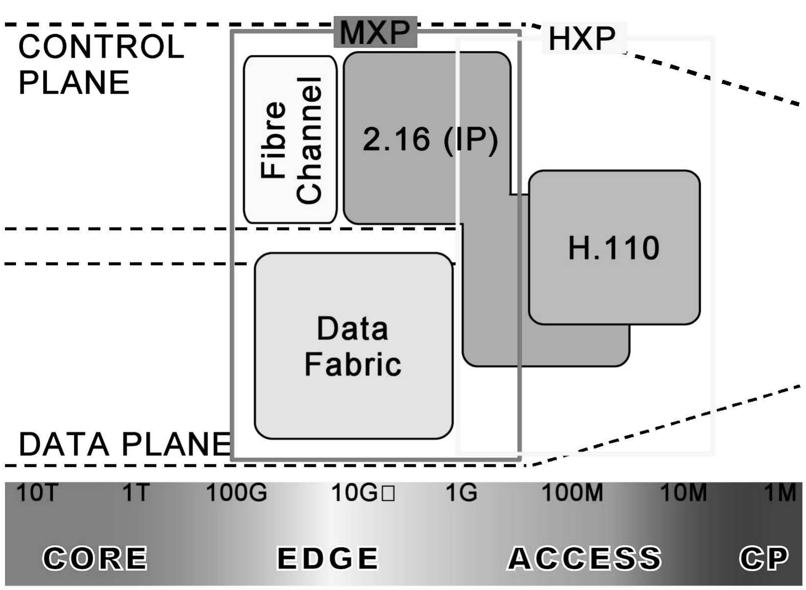

TDM compatibility with a packet transport architecture. One solution to address this need while still leveraging open industry standards is to combine the PICMG 2.16 architecture with PICMG 2.5 (H.110) to create a mixed TDM and IP system. This allows the data traffic to be carried on the IP network and the voice traffic to be carried on the TDM network. In this manner, telecom OEMs are offered a migration strategy that addresses traditional network architectures, while also addressing the needs for packet-based architectures.

Next-Generation Solutions

Motorola Computer Group has addressed the changing needs of the packet world with the introduction of two new architectures, the Multi-Service Packet Transport Platform (MXP) and the H.110 Packet Transport Platform (HXP).

Built to adhere to open industry standards regulated by PICMG, the MXP is a highly flexible platform that can provide fast throughput (more than 700 Gbps) as well as the flexibility and scalability needed to deliver multiple services that can connect to different networks—IP and ATM.

The MXP multiservice architecture takes advantage of the mesh topology to offer a more resilient solution that needs no dependence on a central resource. The mesh architecture, which is being submitted to PICMG for consideration as an open industry standard, enables a more scalable solution than that presented merely by a star network.

While the networks continue to evolve to all packet-based, there remains the need to protect OEM investment in circuit-switched technology. To address this market need, Motorola has augmented its packet transport platform series with the introduction of the HXP. This product provides a clear migration strategy for telecom OEMs by supporting TDM and IP in the same platform.

Comprehensive Packet Solutions for Current and Future Networks

Standards-based systems provide access to compatible products from a broad range of vendors, maximizing the OEM choice. Standards can only be generated from existing technology, however, and in fast-evolving markets, need to be augmented with forward-looking new ideas and technologies.

By developing architectures that adhere to open industry standards regulated by PICMG, while integrating value-added features from Motorola, the telecom OEM is well equipped to meet the needs of a changing marketplace. The packet transport platforms offer the necessary open architectures and speed and flexibility requirements to meet the needs of today's network while providing the foundation to evolve into next-generation solutions.

IP Telephony

Michelle Specktor

Co-Founder and Vice President, Business Development

Gallery IP Telephony, Inc.

Internet protocol (IP) telephony is the emerging communications technology arising from the convergence of the worldwide data infrastructure and the telecommunications networks. The evolution of the public network from a circuit-based infrastructure to a packet-based infrastructure introduces new software-based Class-4 and Class-5 switch alternatives that will enable service providers to reap the benefits of the next-generation network (NGN) IP–telephony revolution with great flexibility in rolling out new revenue-generating opportunities, low overall cost of ownership, platform scalability, and enhanced price-to-performance ratios.

While the IP telephony opportunity is huge and most promising, the mass deployments of softswitches and the transition to an all–packet-based network will take a long time to complete. Before IP telephony becomes a major growth market, the softswitch market needs to mature. Both service providers and vendors face key issues that need to be first resolved, then developed and implemented, before the softswitches and IP–telephony markets could meet their exponential growth forecasts.

The key issues for service providers are demand for selectivity in choosing best-of-breed end-to-end solutions, low overall cost of ownership, enhanced price-to-performance ratios, flexibility and scalability to ever adapt to business growth and new competitive strategies, an increase in quality of service (QoS), and an ability to create new revenue-generating opportunities and support fast rollout of new services.

While the key issues for service providers are in terms of profitability and competitiveness, the equipment and software vendors face a lack of standardization in key areas; hurdles in comprising voice, data, and software development specialists; difficulty in balancing between reliability and scalability; dense interoperability testing and performance testing; financing struggles; and, given the new marketplace circumstances, a more difficult overall sales environment.

The IP–telephony revolution challenge is to combine software advantages with sophisticated telecommunication architecture to create the ultimate highly reliable IP–telephony solution, which in time will replace the traditional circuit switch and will provide a strong economic advantage to service providers offering bundled communication services and enhanced services.

While the traditional voice switching market was dominated by a very small number of vendors (five vendors account for approximately 85 percent of the global public switched telephone network [PSTN] switching market), the vendor environment of IP telephony is vastly different. Historically, each technology revolution has set the stage for new players. Likewise, many new start-up vendors will aim to address the evolving IP–telephony market opportunities, yet, most likely, just a few "right players" meeting the demanding criteria of the teir-1 service providers will be able distinguish themselves as the leading IP-telephony vendors (e.g., Sonus, Gallery IP Telephony, Telcordia Technologies).

A softswitch is a pure software-based switching and control solution that runs on general-purpose industry open platforms to provide the functionality of the traditional time division multiplexing (TDM) switch in a modular, distributed fashion and based on industry standards. The softswitch's key performance criteria are reliability and scalability. It combines the functionality of the traditional Class-4 or Class-5 switch (see the following descriptions), yet its open architecture supports extendibility and scalability, and use of standard complementary components. The softswitch delivers the benefits of cost reduction, service differentiation, and the leveraging of existing networks while migrating to IP–based networks.

The traditional switches are split into two distinct types of switches:

1. Class-5 switches reside at the edge of the PSTN, providing the local exchange/central office (CO) functionality, serving subscribers through local-loop connections, and handling call control as well as attributing special subscriber features and service requirements. Traditionally owned by the local-exchange carriers (LECs) and service providers.
2. Class-4 switches, also known as tandem toll centers, reside at the central part of the network. Essentially, these interface with either other tandem switches or Class-5 switches to complete calls outside of the local area. The role of these tandem switches is simply to control the relay of information from one (geographical) network to another. Traditionally owned by long-distance/international carriers and service providers.

In converged networks, computers are becoming part of the network, and traditional proprietary switches are replaced with standard-based softswitch environments that are composed of software-based ambiguous switches, gateways, application servers, and complementary network components.

The core component of the converged network is called the "call agent," "media gateway controller," or "softswitch."

Class-5 softswitches are infinitely more complex in their design and operation and will account for the majority of added value associated with deploying IP–based networks. "The Class-5 market represents more than 80 percent of the total market opportunity for voice switching and presents the greatest opportunity for services innovation and differentiation."[1]

Although Class-4/tandem softswitches and gateways could be found already deployed by Internet telephony service providers (ITSP) and long-distance/international carriers, generating revenues for players in this niche (e.g., Sonus, ipVerse), they have become somewhat of a commodity. Clearly, it will be the deployments of Class-5 softswitches that will support the service differentiation and costs reductions that next-generation telecommunications companies are seeking, and thus the Class-5 softswitch together with a flexible robust services platform is the ultimate technology required for reaping the greatest benefits of IP telephony.

In summary, these are challenging days for next-generation service providers and carriers, as well as for equipment and software vendors. Never before has there been so much demand for new technologies, and never before has there been so many vendors' offering. But making the right decision has gotten harder. The new IP–telephony infrastructure is not yet fully spread out, and the softswitch technology is not yet mature. However, once the key issues of both the carriers and vendors are resolved, the right vendors with the right technology, partners, strategy, and paths to the market will not only dominate the market, but also drive and accelerate IP–telephony market growth.

Notes

1. Source: *Directions in IP Communications*, CIBC World Markets (November 8, 2001).

VoIP in the Enterprise: Opportunity or Threat for Service Providers?

Laura Thompson
Vice President, Corporate Marketing
Sylantro Systems Corporation

Internet protocol (IP)–private branch exchanges (PBXs) are quickly gaining market-share among recent communications system purchases by enterprises, according to Phillips Infotech and other analyst firms. Large distributed enterprises, such as Household International and Dow Corning, and institutions, such as the University of Arkansas, have made the switch to these newer, IP–based in-house systems. This growing trend toward IP–PBXs presents both an opportunity and a threat to service providers.

The IP–PBX has the potential to further extend and consolidate the powerful hold that customer-premises equipment (CPE) has maintained over enterprises since the mid-1980s. Over the course of 20 years, service-provider–based Centrex offerings were inexorably squeezed out of the enterprise market—and later, the small- to medium-sized–business markets—by feature-rich CPE products that far exceeded the limited functionality, features, and ease-of-use of traditional Centrex.

The IP–PBX trend is just beginning and is likely to follow an easier, faster adoption curve than that of the PBX. Three years from now, according to Gartner, 44 percent of lines shipped will be IP–PBXs if current trends continue.

As a newcomer to the CPE space, the IP–PBX poses an additional, more ominous threat to service-provider revenues and viability. In addition to providing a rich set of features, the IP–PBX also allows businesses with geographically dispersed offices to route interoffice traffic via a data network, consolidating trunks and lines and significantly impacting service providers' toll and long-distance revenues, as well. Thus, the IP–PBX has the potential to effectively blockade service providers from the enterprise market forever in terms of service offerings, increased access lines, and their associated usage revenues.

Enterprise users who purchase IP–PBXs are simply making the best choice that they can among current offerings—and who can blame them? The IP–PBX effect is predicted to continue to grow in strength unless a more compelling service-provider–based offering is provided. A recent study by Research First, a telecommunications market research and consulting firm, found that outside of the small office/home office (SOHO) market, about 90 percent of businesses have installed an in-house phone system. In-house systems such as these are typically replaced every four to seven years. Without a viable service-provider–hosted alternative, much of this opportunity will be lost to the IP–PBX.

With the penetration of on-premises systems at more than 90 percent, traditional Centrex is clearly not perceived by businesses as a viable alternative. In fact, in a study of the Centrex market by Frost & Sullivan, it was found that most businesses view Centrex only as a stopgap measure to use until they can make a decision on an in-house phone system. Other reports have indicated that, due to favorable tariffs, Centrex lines are often less expensive than other access line alternatives—and thus many deployed Centrex lines are in actuality used only for trunking rather than features such as call transfer, park, and pick-up. Traditional Centrex has not and will not be a viable competitor to the IP–PBX.

The Service-Provider Impact

Before the introduction of CPE telephony systems in the early 1980s, service providers were seeing access line growth in the 10 to 12 percent range. After PBXs and key systems became entrenched in the enterprise market, providers' access line growth dropped to 3 to 4 percent in the late 1990s as a result of these on-premises systems' ability to allow multiple users to economically share lines. This decline in growth and revenues has continued due to the rise of other technologies—such as wireless communications, e-mail, and instant messaging—which have diverted even more traffic from service providers' wireline networks. If service providers are to generate new revenues to offset these losses and produce a healthy balance sheet, they must diversify beyond traditional product and service offerings and offer compelling new enterprise-oriented services.

A New Opportunity

Many of the technologies and market forces that have engendered the IP–PBX have also given rise to a new type of products and services that are uniquely designed for service-provider networks and business models. These products, called hosted PBX or IP Centrex, provide a strong alter-

FIGURE 1

New Technologies Impact Access Lines

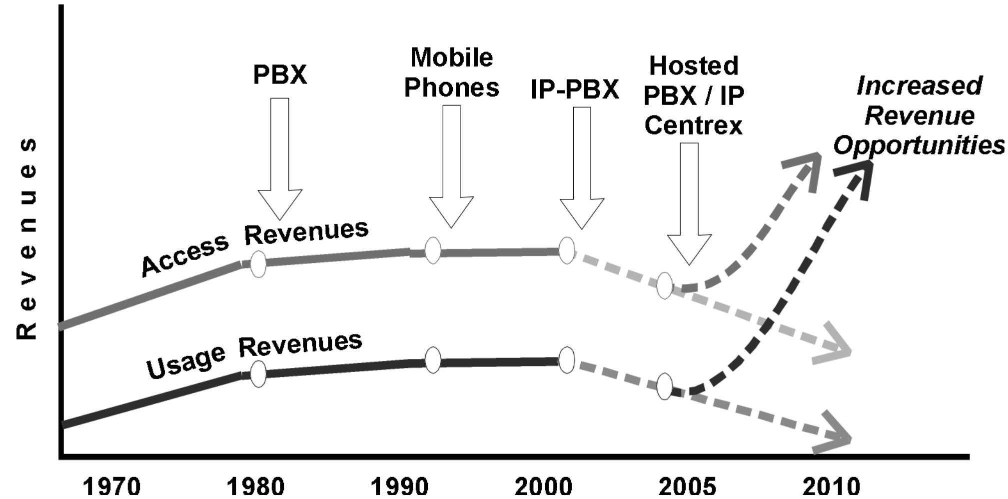

native to customer-premises solutions and can compete effectively against the rising IP–PBX tide. Residing in the service provider's network as application or feature servers, these new products work with existing time division multiplexing (TDM)–based (Class-5) networks and operations support system (OSS) infrastructures and can support a smooth, phased migration path to next-generation softswitch-based networks if that is the ultimate goal.

As the name implies, hosted PBX, or IP Centrex, is based on the same type of deployment model as traditional Centrex. However, by virtue of its IP–based architecture, hosted PBX, or IP Centrex, shares many of the same advantages of IP–PBXs while providing additional benefits that premises-based systems can't offer.

IP–PBXs provide feature delivery, switching, and other functionality over a data network. These systems reside at the customer premises and require internal information technology (IT) staffing resources for day-to-day management, as well as vendor support for major changes. Calls within the premises and interoffice off-premise calls are typically routed over a wide-area network (WAN) or other data network, bypassing the public switched telephone network (PSTN) (and service-provider access charges) entirely. According to Probe Research, up to 70 percent of an enterprise's traffic is inter-employee, especially for large enterprises, so the savings for users—and lost revenues for service providers—quickly become significant.

Hosted PBX, or IP Centrex, is based on a service-provider network-based model, much as traditional Centrex, but offers feature parity as well as cost parity, along with advanced new local and remote management capabilities. With these types of hosted services, all calls, switching, features, and other functionality are distributed to the customer location(s) via broadband services such as T1 or digital subscriber line (xDSL). Just like the IP–PBX, calls are delivered via an internal data network to users within a single premises or over the service provider's network (a virtual WAN for the enterprise) for calls between multiple enterprise locations. However, with IP Centrex services the customer no longer has to invest in or manage a PBX or IP–PBX—they need only to IP–enable their existing business phones with a simple adapter or use new IP phones. No other CPE is needed.

For businesses with multiple sites, interoffice calls can be routed across the service provider's managed data network—retaining these access revenues for the service provider. The business customer benefits, as well. By using these service-provider network resources, business customers gain the advantages of a network without the upfront capital expense of building one or the management headaches and costs associated with maintaining their own network.

End-User Benefits

Hosted solutions such as hosted PBX, or IP Centrex, offer business users the same wide-area dialing advantages and savings as IP–PBXs, without requiring internal IT resources to manage the system. The communications services and the broadband facilities used to deliver hosted PBX capabilities are managed instead by the service provider, simplifying deployment and management for the business customer compared to an on-site IP–PBX. In addition, a service-provider–managed network typically provides a level of voice quality that is indistinguishable from traditional PSTN–based services.

Hosted PBX, or IP Centrex, systems can also meet requirements for telco services such as E-911 and caller ID. These capabilities are identical to, or supersede, those available with traditional TDM–based voice services. Hosted PBX, or IP Centrex, services are offered at a monthly service fee per user, much like traditional Centrex, allowing businesses to avoid large capital outlays of $25,000-plus to upgrade existing PBXs or $75,000-plus to purchase a new on-premises IP–PBX. Through disruptive pricing strategies, service providers can easily compete with, and win against, IP–PBXs.

In addition, hosted PBX and IP Centrex features far exceed those provided by legacy CPE products and easily match those offered by IP–PBXs. The advanced features offered by hosted PBX and IP Centrex are key to driving adoption, because they are far more intuitive than traditional PBX features. Furthermore, the outsourced model employed by IP Centrex reduces both capital expenditures (CAPEX) and operational expenditures (OPEX) in a significant way for business users, which is very attractive in the current economy.

But Will They Outsource?

How willing are businesses to outsource their telephony services? Outsourcing has become a standard practice for many business services such as Web hosting, payroll, and e-mail. Furthermore, the Research First study referenced earlier concluded that business users are extremely frustrated with managing their current phone systems and exhibit a strong interest in outsourcing their telephony services. Outsourcing was viewed not only as a way to reduce the management and maintenance headaches associated with their current CPE phone systems, but also as a means to avoid CAPEX.

Outsourcing allows businesses to focus on their core business, rather than devoting personnel and expense to their telecommunications systems. In the study, focus groups were conducted in several locations across the United States in which generic sample screen shots of a hosted PBX, or IP Centrex, solution were presented. The samples included advanced features such as click-to-call and inbound/outbound/missed call logs, and new browser-based management tools. After viewing the examples, 60 percent of the focus group members expressed a willingness to outsource their telecommunications. Said one survey respondent, "Get me out of the phone business—please!"

In addition, respondents indicated a strong willingness (a "5" on a sale of 1 to 7) to switch to an outsourced model if that was the only way to get some of the advanced features and management capabilities they were shown. All of this explains why outsourcing solutions such as IP Centrex are being well received in the market—and post-study field experiences have reaffirmed the research project's results.

The IP Centrex Advantage

Technology providers such as Sylantro Systems have used the flexibility of IP to incorporate customer-pleasing new applications into their hosted PBX, or IP Centrex, products. Compelling features such as popular mobile phone-type features including click-to-call and call logs, Microsoft Outlook integration, and hosted directories proved very desirable to business users in the Research First study and in deployments.

These new hosted PBX, or IP Centrex, systems support newer IP phones or IP–enabled legacy phones and offer user-friendly features via new uses of the phones' liquid crystal displays (LCDs) and soft keys. Other user interfaces are offered via browsers Microsoft Outlook, or WAP–enabled mobile phones. Customizable call treatments allow users to set up find me/follow me scenarios for different types of callers, such as VIPs, family, or friends. A call queuing and distribution option helps businesses manage peak call volumes on a temporary or permanent basis—a much more affordable option to buying and managing an in-house automatic call distributor (ACD) system.

While the service provider maintains the servers, some hosted PBX, or IP Centrex, solutions offer a rich browser-based portal that allows office receptionists or managers to perform moves, adds, and changes (MACs) quickly and easily. The Yankee Group estimates that businesses spend almost $6 billion per year on phone-system maintenance, including MACs. Through portals such as these, businesses can largely eliminate these expensive outlays and gain control over MACs themselves—without the frustration of depending on outside service contractors.

These routine tasks have historically been a major source of frustration, delays, and expense for users of premises-based telephone systems. The self-administration capabilities offered by hosted PBX and IP Centrex solutions are attractive to business users because the capabilities allow fast responses to changing business needs and conditions without a lengthy wait for an expensive service call (typically a minimum $200 charge per call). Furthermore, the portal allows administrative staff—rather than expensive IT staff—to manage routine changes.

Browser-based portals also allow end-users to easily access, manage, and control their communications options on the fly, without star codes or other complicated programming gymnastics. These portal interfaces help users, for the first time, to take full advantage of communications features that enhance productivity and increase accessibility and flexibility.

Unlike premises-based telephone systems, hosted solutions are easily upgraded. Service providers can allow customers to change service packages at any time they choose and to upsize or downsize the services as needed or as their businesses flex in number of employees or services needed. Premises-based solutions such as IP–PBXs simply don't offer these same advantages. The ability to "right-size" communications services to fit changing needs is critical in today's economy and allows businesses to use only the resources they need, without incurring unneeded expenses.

Building New Revenues

During the next few years, service providers have a unique opportunity to generate much-needed new revenues and recover some of the business that they lost to premises-based solutions over the last two decades. Through the added flexibility afforded by hosted PBX and IP Centrex, service providers have a powerful set of tools that will allow them to easily compete with and outsell IP–PBX offerings.

Many of these advantages are a function of the service provider's market presence. For example, the study by Research First found that the first company most businesses think of when considering a new telephone system is their phone company. Service providers have built extensive brand recognition and customer awareness that can quickly translate to sales opportunities for a compelling new service offering such as IP Centrex.

Most service providers also have a knowledgeable sales force in place that is already selling traditional PBXs and Centrex and that can easily transition to selling new services such as IP Centrex. Some innovative service providers with IP Centrex offerings are signing up former traditional Centrex resellers to market these services, with great results.

Interestingly, field experience has shown that, while hosted PBX and IP Centrex solutions might be useful for migrating Centrex users to IP–based telephony offerings, the biggest opportunity is in replacing traditional PBXs and key systems. This market segment is a true incremental opportunity for service providers.

Summary

For all of these reasons and many more, IP Centrex has emerged as a viable, competitive alternative to IP–PBXs and other CPE solutions. The innovative service providers offering hosted PBX and IP Centrex today are winning sales away from IP–PBXs. Oftentimes, end-users are replacing their phone system much sooner than the average (typically four to seven years), since there is no CAPEX penalty for replacement, and they gain access to the extensive features and easy manageability offered by these hosted services. Through hosted PBX and IP Centrex, users receive all of the new call-management functionality that IP–PBXs offer—and new usability and control that traditional PBXs never offered—all without the up-front expense of a premises-based solution.

For the first time in years, Centrex—in its IP–based form—is fully competitive with CPE solutions, and where businesses have the choice, these new services are winning hands down against IP–PBXs. Hosted PBX and IP Centrex solutions are proving themselves in the market and have amassed solid evidence that they can combat the IP–PBX trend and help reverse the decline in access lines and revenues for service providers. IP Centrex has come into its own, but the window of opportunity, before the IP–PBX reaches critical mass, is now.

Requirements for the Next-Generation Intelligent Data Network

Eric Von Gonten
Senior Internet Architect, Internet Architecture
Genuity

This paper will discuss requirements for the next-generation intelligent data network, providing an overall view of the network, including backbone and metropolitan areas. This paper does not pertain to Genuity's network in particular, but rather networks in general.

Backbone Network Architectures

Current backbone network architectures primarily consist of 10 Gigabit per second (Gbps) dense wavelength division multiplexing (DWDM) transport systems that were deployed between 1997 and 1999. These networks were built with passive, point-to-point (PTP) DWDM technology with anywhere from eight to 32 lambdas of transport capability. In addition, these networks were overlaid with state-of-the-art 10 Gbps synchronous optical network (SONET) four-fiber bidirectional line-switched rings (BLSRs) to ensure reliable architectures for service providers and their customers.

Cost Concerns

Four-fiber BLSR architectures are highly reliability. However, they are very costly, include many complicated components, are difficult to provision, and require long build times. Significant capital expenditures (CAPEX) must be made to purchase the equipment necessary to support this level of reliability. Four-fiber rings are designed so that half the bandwidth is dedicated to working traffic and the other half is dedicated to protection, and the architecture requires the deployment of many redundant transmitters. All of this adds up to a significant cost disadvantage.

Revenue Concerns

Under four-fiber BLSR architectures, traffic can be provisioned on only the working path of the ring, while the protection path sits idle, without generating revenue for the carrier. If service providers had access to the unused protection bandwidth, the network efficiency would be improved. Without the ability to provision on the protection path, service providers have spent large amounts of money for optical components to build protection, without generating additional revenue.

In a network where multiple layers of rings are deployed across a large geographic area, provisioning traffic across numerous rings becomes quite challenging. Because rings can be anywhere from hundreds to thousands of kilometers in circumference and must be completely built out before they can be used, build times are generally long. The business issues associated with four-fiber BLSR rings are cost, excessive equipment, ineffective use of bandwidth, complex provisioning, and long build times.

Reliability Concerns

There are also concerns about the reliability of four-fiber ring systems due to the inherent operational complexity. The DWDM transport portion of the network tends to be very reliable, as it is primarily composed of amplifiers and passive filters. DWDM systems of the past required signals to be de-multiplexed, regenerated, and multiplexed every few hundred kilometers. On a large network, this resulted in many fiber patch cords and connections between the DWDM layer and the active transmitters of the SONET four-fiber rings. Additionally, this type of regeneration required service providers to install SONET bays and numerous optical-to-electrical-to-optical (O–E–O) cards to regenerate both the working and protection channels of the many already deployed four-fiber rings. The requirement for so many patch cords, connections, and electrical interfaces increases the opportunity for an outage to occur either by equipment failure or user error, which can impact the reliability of network.

Architectures of the past used Internet protocol (IP) over asynchronous transfer mode (ATM) over SONET over DWDM. This architecture utilized four different network layers to offer services to customers. The more layers that were deployed resulted in more required components for providers to purchase and provided more opportunity for failures due to the large number of components installed. As technology continues to advance, new transport mechanisms and architectures will arise and will reduce some of these concerns.

Routers Plus Photonics Architecture

Other architectures under current consideration by the industry use only two layers: an optical layer and a service or IP layer. *Figure 1* is an end-to-end view of the service layer of the future resulting in Layer-1 to Layer-3 convergence. The vision here is to take 10 Gbps IP traffic directly over route-diverse optical wavelengths bypassing both the ATM and SONET equipment and relying on the optical layer and IP routers to provide the protection. With the recent development of 10 Gbps interfaces on routers and with optical layer support for 10 Gbps transport, the requirement to use SONET equipment to aggregate lower-speed signals up to a 10 Gbps signal is no longer necessary. This architecture will continue to evolve and use high-speed terabit routers, which can have dozens to hundreds of interfaces with speeds from 10 Gbps to 40 Gbps and faster on each interface. To support this architecture, the network must use a highly scalable optical layer that supports terabits of capacity and has the ability to provide bandwidth on demand.

The use of ultra-long-haul (ULH) DWDM technology to provide city-to-city connectivity can be used for this type of architecture with improved overall network reliability. In the past, a signal sent from Los Angeles to Miami would require electrical regeneration at many locations. However, with an ULH DWDM system, little or no intermediate electrical regeneration would be required. Therefore, ULH DWDM technology may actually increase the reliability of the network by reducing the number of patch cord, connections, electrical interfaces, and intermediate regeneration sites compared to previous DWDM transport systems.

The use of ULH technologies to transport IP directly over wavelengths is the first step toward the next-generation intelligent optical network. The benefits associated with this architecture include a cost-effective transport, reduction in O–E–O requirements, higher reliability through reduction of components, operational simplification resulting from bypassing higher-order network layers, and improved provisioning intervals.

Metropolitan Network Architectures

Current Metropolitan Network Architecture

The current industry-standard metropolitan network includes fiber, limited DWDM, SONET, and a service-aggregation component. There are many different hardware platforms, components, and software platforms that coexist in today's metropolitan networks. All these complexities result in many opportunities for inefficiency (see *Figure 2*).

Reliability Concerns

The more components that exist in a network, the higher the possibility that a failure may occur. In addition to multiple hardware and software platforms, a complex metropolitan architecture consists of many patch cords, electrical cables, connections, and interfaces that increase the opportunity for an outage due to equipment failure or user error, thus impacting the reliability of the metropolitan network

Hardware is not the only item that can affect the reliability of a network. Software upgrades have many complexities; time and careful planning must be taken when they are performed. Over the years, there have been more than a few carriers that have had outages associated with software upgrades. Outages could be associated with people following incorrect procedures, an upgrade of one system not interoperating properly with another system, or even a bug in the software. Regardless of the problem, when carriers experience outages, customers and revenue are potentially impacted.

FIGURE 1

Routers Plus Photonics Architecture

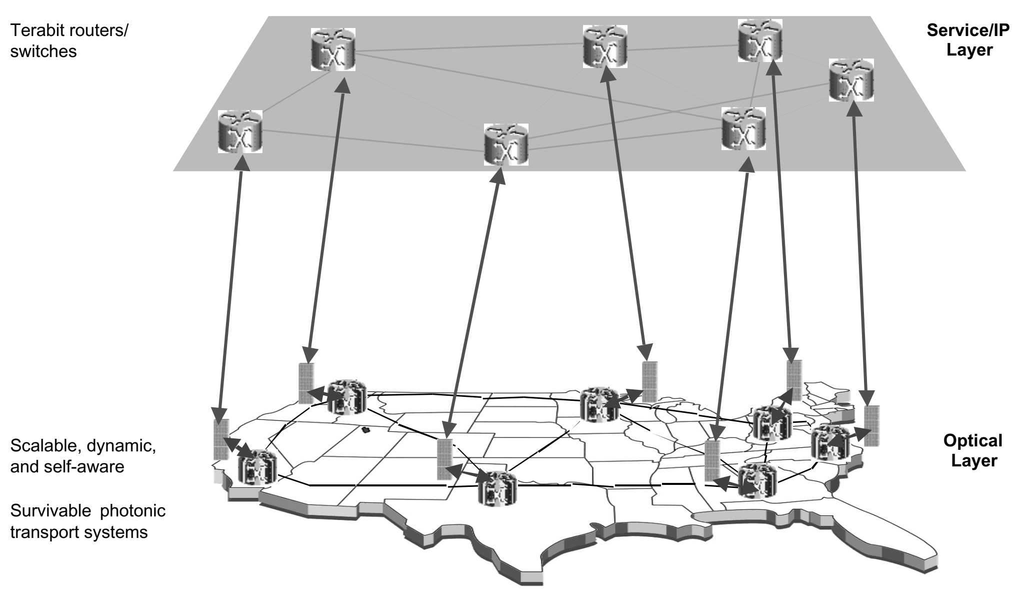

FIGURE 2

Metropolitan Network Migration

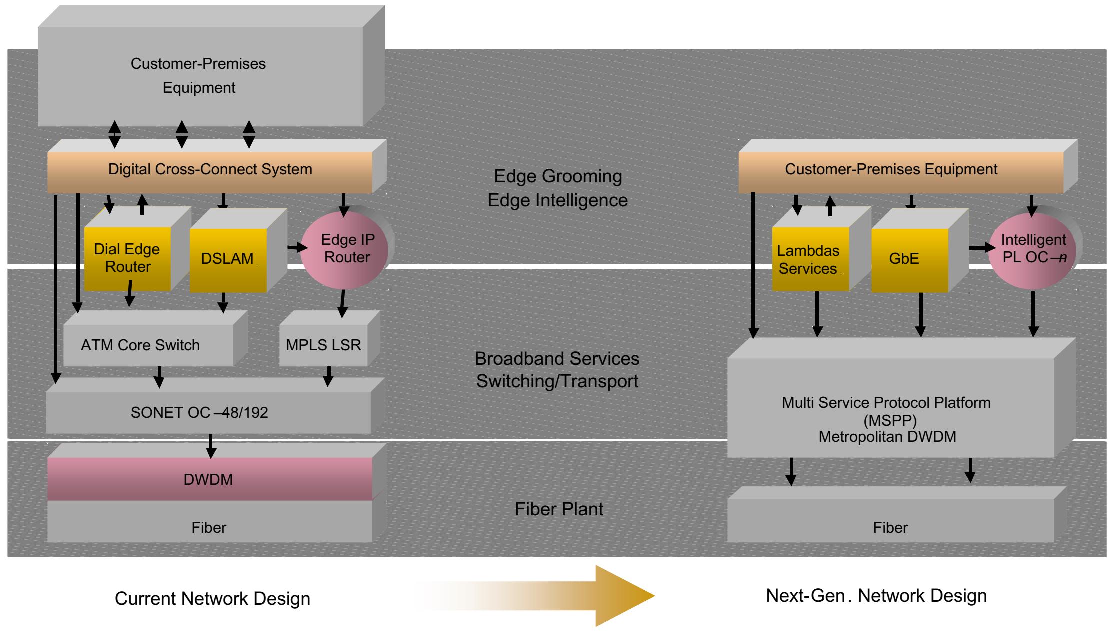

The more components in the network and the more software upgrades you perform result in a higher risk of a network outage. Current complex metropolitan architectures, with many components and management systems, present operators with many opportunities to make slight mistakes that could cause network outages.

Next-Generation Metropolitan Network Architecture

The next-generation metropolitan network will be a consolidated network where a single integrated access device (IAD) or multiservice protocol platform (MSPP) will replace many of the components of previous architectures. These MSPPs will offer a variety of service interfaces such as SONET, IP, Ethernet, ATM, etc., along with test-access in one platform resulting in simplified network architecture. MSPPs seem to be very popular, although the operational aspects of this approach present some challenges that need to be considered in detail.

In the metropolitan environment, we must consider churn. When a customer requests a service upgrade (optical carrier [OC]–3 to an OC–12) there are many things that occur. A technician is required to drive out to a site, install a new interface card, connect patch cords from the interface to a demarcation patch panel, coordinate the interconnection with the customers technician, provision the circuit, and finally test the circuit. All of this takes time and money. While this process is taking place, the customer is not without service, as they are still utilizing the original OC–3 circuit. When the customer has accepted the new circuit, the old OC–3 circuit must be decommissioned through a process that includes removing the old OC–3 interface card, removing the patch cords from the card to the patch panel, and removing the provisioning information from the database of record. This manually intensive process of turning up new circuits and decommissioning old circuits presents many opportunities for operator error, which may result in a customer outage.

Providers can minimize opportunities for outages and simultaneously offer more cost-effective services to customers by immediately providing the most cost-effective, highest bit-rate interface possible. This could be a Gigabit Ethernet connection or a single interface card that supports multiple bit rates, and service upgrades to the customer can be completed by remotely increasing the data rate of the multirate card. This method lessens the field work associated with ordering and installing new cards, running new patch ports, removing old cards, and removing old patch cords—ultimately reducing the possibility of errors that could degrade network reliability. In addition, this eliminates the cost of purchasing new cards for upgrades and reduces the need for old cards to be returned to the carrier's inventory.

The Metropolitan Network Is about End-to-End Services

Figure 3 shows the technology of the metropolitan network. The IP backbone is on the right, the metropolitan aggregation and distribution is in the middle, and the last mile (thousands of subscribers with customer-premises equipment [CPE]) is on the left.

An additional driver for this MSPP technology is the huge demand for reliable end-to-end services offering broadband delivery of content and services. Services that were not even envisioned only three to four years ago are now a reality. Some examples include broadcast video over IP, voice over IP (VoIP), interactive gaming, collaborative services, and the huge growth in time-sensitive traffic, such as on-line trading. MSPPs also provide a way to aggregate many different

FIGURE 3

The Metropolitan Network Is About End-to-End Services

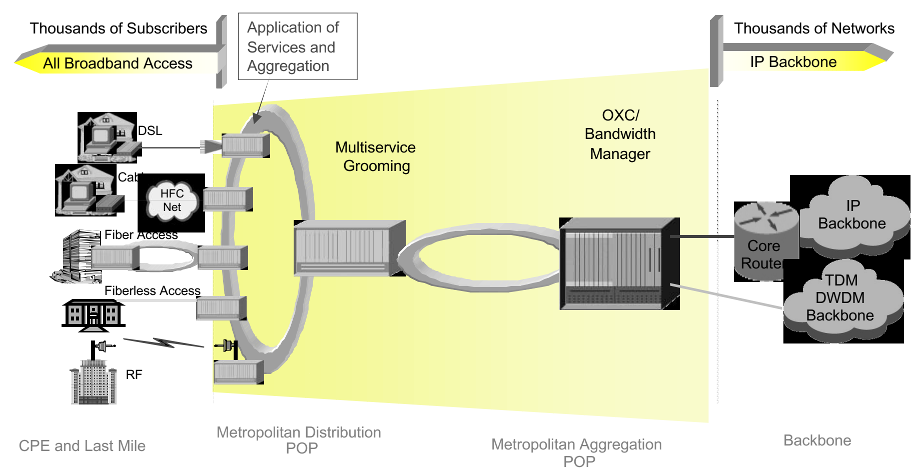

types of traffic into high-speed connections (data or TDM) that can easily be handed off to the backbone network.

When providing an end-to-end service for customers, the reliability of the service depends on the combined architectures of the metropolitan and backbone networks.

Standards and Protocols

Networks of the future must be able to provision bandwidth dynamically through the network based on machine-to-machine signaling. This is just beginning to happen due to Layer-3and Layer-1 convergence. However, it will not work without standards-based protocols. Guidance for building these networks of the future can be found in the legacy telco networks' signaling system 7 (SS7).

Signaling and routing protocols that will enable better management of the network are currently under development. Some key protection and recovery protocols that are currently used are open shortest path first (OSPF), intermediate system–to–intermediate system (IS–IS), border gateway protocol (BGP), and service-location protocol (SLP).

Some other emerging standards and protocols are generalized multiprotocol label switching (GMPLS), optical user-to-network interfaces (O–UNIs), optical network-to-network interfaces (O-NNIs), optical label-distribution protocol (OLDP), and automated switched transport network (ASTN). Some of the forums currently working in this area consist of the Optical Domain Service Interconnect (ODSI), Optical Internetworking Forum (OIF), Internet Engineering Task Force (IETF), and International Telecommunication Union (ITU).

These standards will enable service providers to provision and manage services in multivendor environments. A large network can consist of a very large backbone built with equipment from two or three different vendors connected to numerous metropolitan networks that may have equipment from five to 10 different vendors. With one vendor in one portion of the metropolitan network, another vendor in the backbone, and even a third vendor in another metropolitan portion of the network, signaling and routing protocols and standards are required to make it possible to manage and dynamically provision end-to-end circuits across the network.

End-to-End Network-Management Capability

Today, only carriers that have end-to-end visibility of both the metropolitan and backbone networks can easily accomplish customer network management (CNM) for end-to-end service delivery. CNM and bandwidth-on-demand services will enable over-subscription of the network, thus increasing revenues for service providers. CNM will also provide customers with instantaneous capabilities, such as performance monitoring characteristics and the circuit reliability.

CNM and future signaling and routing protocols (e.g., GMPLS) will enable bandwidth to be provisioned dynamically through the network based on machine-to-machine signaling. This will allow customers to pay for the network as they use it (dynamically), rather than as it is manually provisioned. Signaling and routing protocols will allow bandwidth to be provisioned dynamically and allow CNM capabilities over a multivendor and/or multicarrier network. This end-to-end network must be operationally friendly, accommodate service-level agreements (SLAs), and be able to scale with the network.

Services and Performance Characteristics of the Future Network

Intelligent software and standard protocols that enable bandwidth on demand are also the primary developments that will enable many new protection mechanisms and classes of service (CoS).

Premium Service: Dedicated Mesh for Near Real-Time Traffic Restoration

The premium service is the real-time traffic restoration (< 50 ms). This service level is appropriate for traffic that is intolerant of delays, such as voice, VoIP, video broadcasting, interactive gaming, and collaborative services.

Today, premium service is accomplished using SONET four-fiber BLSRs. As previously indicated, large amounts of equipment are required to provide this level of service using traditional four-fiber BLSR systems. The next-generation routed photonic network will enable mesh architecture and near real-time traffic restoration to be accomplished by use of dedicated 1+1 optical protection, similar to a tail-end switch or hot standby capability.

The advantage of using a mesh to provide near real-time traffic restoration is the ability to provide higher utilization of network resources than the four-fiber BLSR system. As a result, other services cannot be offered throughout the mesh by utilizing the protection bandwidth.

Gold Service: Predefined Traffic Protection

Gold service uses mesh protection to provide longer restoration times (300 ms to five seconds), and gold service traffic uses predefined protection in the mesh. This service level is appropriate for traffic that can tolerate some delay within defined SLA boundaries. This level provides M-for-N mesh-protected service with no degradation in the CoS. M-for-N protection with guaranteed priority and predetermined routes provides greater reliability, faster restoration, and a deterministic behavior. M-for-N protection with best-available routes, strictly determined by the state of the network, provides recovery that can still be sub-second but that may exceed that based on the state of the network. Throttled-down service at the optical layer ensures that the data layer does not have to deal with loss of adjacencies and associated issues.

Silver Service: Best-Effort Traffic Restoration

Silver service is enabled from a network that searches tables for the best route of recovery. This service may tolerate delays in the seconds to minutes range. Applications include e-mail, store-and-forward traffic applications, storage-area networks (SANs), and IP over DWDM (where IP provides the protection).

This service can be sold to carriers who want to establish IP over wavelength connections on different carriers networks with diverse fiber routes from their own, thus offering Layer-3 protection and physical network topology diversity.

Bronze Service: Unprotected Traffic

Bronze service has no recovery mechanism. Applications include IP over this service and swapping bandwidth with other carriers to provide network topology diversity (different transport networks with potentially different fiber routes) where another layer of provides the protection (Layer 3).

Lead Service: Pre-Emptible Traffic

Pre-emptible traffic is the lowest level of service available. Like unprotected traffic, this service provides no recovery. Here, traffic may actually be dropped to provide protection to other traffic. This service is generally up and running but in the event of a network outage, will be dropped as other traffic uses this capacity for protection. Applications include swapping bandwidth or wavelengths with other carriers for protection and low-priority locally stored data. The benefit of this service for the carrier is the ability to generate new sources of revenue off of the existing operational equipment.

Next-Generation Backbone and Metropolitan Network Architectures Combined: Routed Photonic Network

Signaling and routing standards enable a next-generation network that provides intelligence and automated capabilities. The network, which starts with the physical fiber infrastructure, manages traffic in a routed photonic network. A DWDM layer with both long-haul and ultra-long-haul capabilities is placed on the fiber, providing virtual wavelengths or optical virtual private networks (VPNs). Optical switches are then placed at junction locations to provide switching and routing capabilities. Then there is the multiservice distribution, or metropolitan network (see *Figure 4*), which aggregates services from the edge to be transported on the backbone.

FIGURE 4

Routed Photonic Network

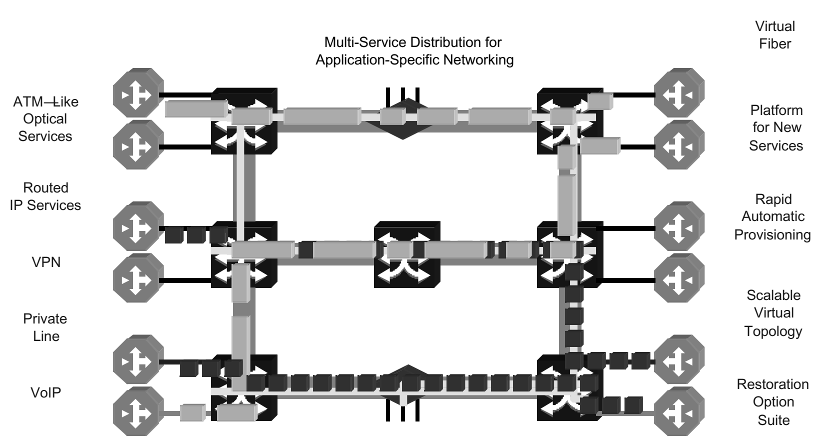

Software, signaling and routing protocols, and standards provide the network with intelligence, which allows the metropolitan and backbone networks to be self-aware of their resources and to operate together to dynamically allocate bandwidth, offer CNM capabilities, and provide end-to-end services to each customer based on the customers individual SLA.

All of these components will enable the future routed photonic network to exist. Many new protection mechanisms and service offerings will arise through this network and will benefit service providers and customers.

Making Next-Generation Networks Possible

Ultimately next-generation networks must deliver end-to-end services to customers cost-effectively. They must be scalable, flexible, and provide for end-to-end network management.

To achieve cost-effective next-generation networks, the telecommunications industry (service providers and vendors) must work together to succeed in delivering business solutions that meet and exceed the constantly changing expectations of a bandwidth-intensive, service-focused customer base.

Parlay API Specifications: An Overview

Avni Wala
Software Engineer, Telecom Applications Group
Future Software Limited, India

Abstract

A number of groups are currently developing technologies aimed at evolving and enhancing the capabilities of intelligent networks (INs). This paper describes one such technology developed by the Parlay consortium, namely Parlay API. This paper describes the Parlay application programming interface (API) specification, its architecture, its advantages, and its place in the current scenario, as well as means of deploying APIs and an example illustrating the use of APIs.

Introduction

Recently, immense exposure to the Internet has raised customers' expectations of the service features that should be offered by the public telecommunications infrastructure. The users expect that the services should support a mix of media types, allow easy customization, and be available on demand, regardless of the location or the capabilities of their terminal equipment. To meet these demands in a cost-effective manner and to survive in a volatile marketplace, the challenge is to deploy the technological components required to realize these services. As the level of interconnection between fixed, mobile, Internet, and enterprise networks increases, a key component in ensuring these services will be the availability of a common platform for their development and delivery. Of course, the platform should leverage the existing systems as much as possible—to enhance the service-delivery capabilities, rather than totally replace them. One such platform, specified by a Parlay industry working group, aims to provide an object-oriented service-control API that is independent of underlying communications technologies (e.g., public switch telephone networks [PSTNs] as well as wireless and Internet protocol [IP] networks). The APIs are specified in unified modeling language (UML) and are designed to support all major middleware technologies (e.g., distributed component object model [DCOM], common object request broker architecture [CORBA], Java platform, etc.)

What is Parlay Group?

The Parlay Group is a nonprofit entity. It is incorporated in California and managed by Enterprise Ventures.

Parlay Members

The following companies are a part of the Parlay consortium:

- Aepona
- Alcatel
- AT&T
- BT
- Cisco Systems
- Ericsson
- France Telecom
- Fujitsu
- High Definition Systems, AB
- IBM
- Lucent Technologies
- Microsoft
- NEC Corporation
- NET4CALL
- Nokia
- SBC–TRI
- Siemens AG
- SS8 Networks, Inc.
- Tundo Communications
- Ulticom, Inc.
- Westwave Communications

Aims of the Parlay Group

The Parlay Group has generated an API specification providing access to network information and control of a range of network capabilities. The API is specified in UML and is designed to support all major middleware technologies. The Parlay Group aims to specify an object-oriented service-control API that is independent of underlying communications technologies to encourage the widest possible range of market players to develop and offer advanced telecommunications services.

Current Scenario

Currently, intelligent capabilities in conventional switched networks are exclusively under the control of network operators because of incompatible standards, network security, and integrity. Hence, we have two scenarios:

FIGURE 1

1. Network-centric communications-service delivery mechanisms (e.g., advanced intelligent networks [AINs], INs, etc.), which run "inside" the network domain. Such a mechanism is easily manageable and robust, and hence is good for simple mass-market applications. But the problem with such a mechanism is that it cannot access data in the enterprise domain for domain-specific decision-making.
2. Edge-of-network service delivery mechanisms (e.g., enterprise computer telephony integration [CTI]), which run "outside" the network domain. Such a mechanism is good for creating services that meet the specific needs of a customer but generally cannot access critical information and capabilities within the core network.

Figure 1 shows a typical network-centric scenario:

- The top layer represents applications.
- The middle layer represents service components that are reused by many applications.
- The lower layer represents different network technologies.
- All the three tiers of this architecture are owned and operated exclusively by network operators.

Disadvantages
The major disadvantage is the difficulty in achieving the flexibility to deploy many customized versions of services to different customer groups, as the network operator is responsible for the creation and operation of all applications. Also, application development and integration testing take a long time, thus the time to market for new applications is long.

Introduction of Parlay

The advantages of both network-centric and edge-of-network mechanisms have been incorporated in Parlay specifications. Parlay is an umbrella architecture that provides network independence and application portability. Parlay APIs enable a new generation of off-the-shelf network applications and components (e.g., messaging, mobility, end-to-end quality of service [QoS], etc.) to be developed by application providers (e.g., independent software vendors [ISVs], application service providers [ASPs]) independent of the underlying voice/multimedia network. Parlay APIs are platform/vendor/technology independent to ensure the implementation of multiple technologies. The Parlay Group, working together with standards bodies, directs the industry to design and develop new services and allows those services to actually drive information technology (IT) networks.

Where Does Parlay Fit into the Existing Architecture?

As shown in *Figure 2*, the Parlay APIs can be implemented on top of the middleware, which can be any implementation of CORBA, DCOM, or the Java platform. Some of interfaces will be implemented on the client side (i.e., the application side) to collect the responses given by the interfaces that are implemented on the server side (i.e., the network provider's side). The Parlay server-side interfaces will collect the information from the underlying network, and then will present the required information to the client application using the call-back mechanism.

How Can the API Be Deployed?

The main specification is defined in technology-independent UML. Deployment of the API can be via technologies such as DCOM, CORBA, or Java. Technology-dependent interface definition language (IDL) specifications in Microsoft IDL and CORBA IDL also exist. APIs can be implemented using any of the languages that support object-oriented characteristics.

A Conceptual View of Parlay APIs

As illustrated in *Figure 3*, the API resides between the application layer and the service-component layer. By enabling access to network capabilities via an API, service providers,

FIGURE 2

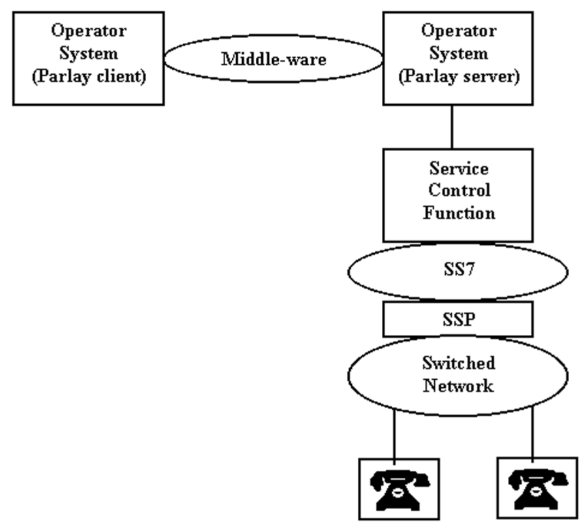

ISVs, and other developers in the IT and telecommunications industries are empowered to generate a new range of applications that benefit from, and add value to, functionality resident in public and private communications networks. These services can then be used within the enterprise or are sold to other enterprises.

The Parlay API Architecture

As illustrated in *Figure 4*, the Parlay APIs consists of two categories of interfaces:

1. Framework interfaces
2. Service interfaces

Framework Interfaces
Framework interfaces provide supporting capabilities necessary for the service interfaces to be open, secure, and manageable. The framework can be considered as a number of functional building blocks and is independent of any of the Parlay services.

Functions provided by framework interfaces include the following:

- Service registration, subscription, and discovery
- Authentication and authorization
- Integrity management

Thus, the framework should authenticate the users of the API, prevent non-authenticated users from accessing both the service and the framework interfaces, and provide the users with means to ensure the integrity of their transmissions.

The framework interface can be divided into the following categories:

1. ***Trust and Security:*** It enables the application to contact the framework, mutually authenticate, and sign for use of services.
2. ***Discovery:*** It supports properly qualified searches for services available to the application.
3. ***Event Notification:*** It allows asynchronous event passing between the application and the gateway.
4. ***Integrity Management:*** It performs various functions to monitor and control the operations of the gateway and application, and assists in recovery from faults.
5. ***Subscription Management:*** It allows enterprise applications to explicitly subscribe to new services and makes them available for use.

Benefits of framework interfaces:

- Common security, service location, and registration framework to which application developers can write. These applications would use the services presented by the Parlay API.
- Provides service brokering, security, billing, and operation support and maintenance (OS&M) capabilities consistently across any Parlay-enabled framework.

Service Interfaces
Service interfaces contain mechanisms by which applications access underlying network capabilities and information.

Functions provided by service interfaces include the following:

- Allowing access to traditional network capabilities, such as call management, messaging, and user interaction.
- The service interfaces also include generic application interfaces to ease the deployment of communications applications.

The service interfaces are divided into five categories:

1. Call-control service
2. User interaction service
3. Messaging service
4. Connectivity manager
5. Mobility service

FIGURE 3

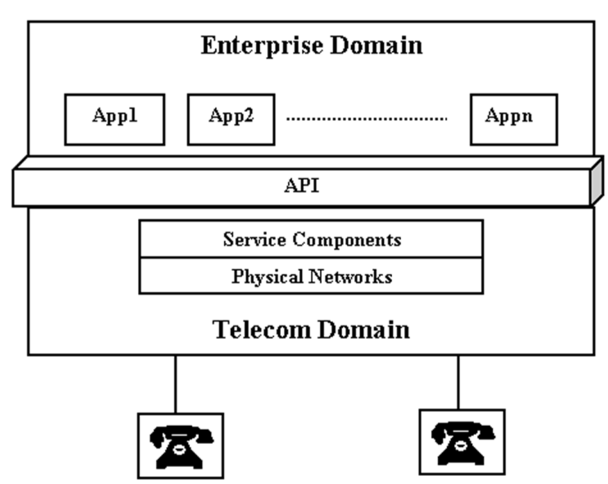

FIGURE 4

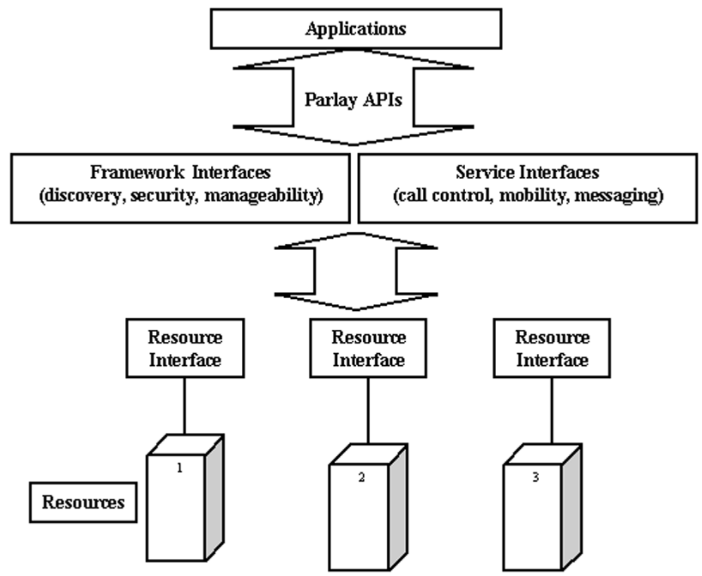

Call-Control Service
This API creates applications that use centralized call routing for services that span PSTN, mobile, and voice over IP (VoIP) networks. It has tiered capabilities from very basic call control to multiparty multimedia, so that applications use only the appropriate complexity.

The call control can be a simple call control, multimedia, or conferencing.

Benefits of call-control service APIs include the following:

- Multimedia applications can signal across all network-specific protocol mediums, hiding underlying protocols.
- The APIs enable multimedia services development that span multiple networks.

User Interaction Service
User interaction service, when used in conjunction with call control, enables applications to control interactive voice systems, specify prompts, and collect user input. It is used to send and receive information to and from the user. Any responses that will go from the application to the user, and vice versa, will make use of the user interaction service APIs.

Messaging Service
It has capabilities to send, retrieve, and store messages. It explores limited functionality to store-and-forward messaging systems for voice and e-mail messaging.

Benefits of messaging service APIs:

- They simplify developers' abilities to write applications for messaging that integrate other services.
- Service providers can offer cross-service messaging applications and ensure application portability.

Connectivity Manager
A connectivity manager enables applications to provision QoS–specified pipes across IP networks. It will allow an application to be developed to configure virtual circuits (VCs) and selectable QoS levels on an IP–based network.

Benefits of a connectivity manager include the following:

- It enables user-configurable QoS multiservice applications across IP networks.
- Configurable QoS levels on IP backbone can be integrated with other applications.

Mobility Service
Mobility service presents mobile user-location information through APsI to be used to develop applications for network services. Using the mobility-service APIs, applications can obtain the geographical location and status of fixed, mobile, or IP–based telephony users. Also, the location information based on the network-related information, rather than geographical coordinates, can be retrieved via these APIs.

Benefits of mobility service include the following:

- It provides APIs for the location of mobile users for emergency, geographical, and other applications.
- Mobile-user location information adds unique functionality to various applications.

An Example Using Parlay APIs

A company, ABC Repairs, offers a product support service for its customers. The software system that it uses comprises a workflow engine that integrates a fault-logging system, home-based diagnostic experts, and a mobile repair fleet.

A typical workflow for the support service is shown in *Figure 5*, and the states through which it will go are numbered from 1 to 12. The descriptions of the states are given as a business scenario:

The business scenario:

1. A customer phones the 0800 support line.
2. The call-control API sends a message to the application identifying the calling-line identity (CLI).
3. If the CLI is not recognized as belonging to a customer, the application instructs the call-control API to route the call to the interactive voice response (IVR) system.
4. The prompts played to the customer and the responses are all passed across the user-interaction Parlay API.
5. When the customer and site location are recognized, the application uses its customer database to ascertain the product type and the level of service for which the customer has paid.
6. If the service call is to be dealt with immediately, the application releases an IVR and routes the call using the call-control API to an available agent qualified for the particular product.
7. The agent takes the customer through a basic diagnosis system and logs the report into the workflow databases.
8. The workflow engine identifies an appropriate mobile repair technician to send to the customer site, using the staff skills database.
9. The mobility APIs are used to locate the best-fit technician for the job, and the completion time for the technician's current job is estimated.
10. The agent receives the technician's estimated arrival time and name.
11. The customer is informed of the same.
12. The mobile technician is informed and the fault report is downloaded to the technician's computer.
13. The calls that the employees are unable to answer are diverted to the voice-mail system.
14. When the caller leaves a message, the messaging API will alert an application that may send an alerting e-mail to the staff member.
15. The connectivity manager API is used by the training and conferencing applications to schedule guaranteed QoS pipes between the distributed offices for the duration of high-bandwidth, real-time conferences.

Conclusions

The focus of the Parlay Group is on producing an API specification that enables enterprises outside of the network operator's domain to access network information and control a range of network capabilities. The Parlay Group believes that there is a tremendous growth opportunity for network operators, service providers, and application developers through the adoption of the Parlay API. Customers will benefit as business needs are met more cost-effectively and quickly using the capabilities offered via the Parlay API.

The following is a summary of the major goals of the Parlay Group:

- Define an API specification that provides application developers a consistent, network-independent view of functionality that they can employ.
- Ensure that the API specification provides security and integrity for the network as well as enterprise domains.
- Encourage IT and telecom industries to implement and deploy the Parlay API globally.

FIGURE 5

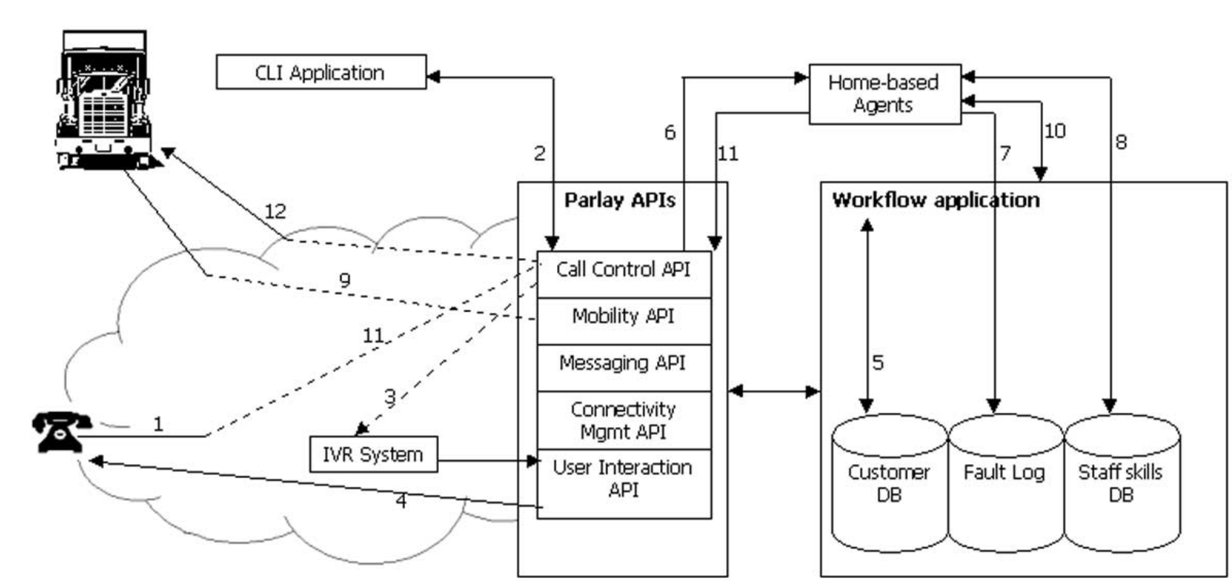

- Encourage application developers to create new and innovative applications using the APIs to solve customers' needs.
- Work with the IT industry to encourage the production of software development kits (SDKs) that make the process of developing applications exploiting the Parlay API straightforward and efficient.

References

1. Parlay API Specification 1.2 – Core Specification document, www.parlay.org
2. Parlay Group Press Release, May 1998, www.parlay.org
3. Parlay Group Press Release, November 1998, www.parlay.org
4. Guaranteed Call-Delivery Demonstration, 1998, www.parlay.org
5. Parlay Group Press Release, June 1999, www.parlay.org
6. UML Version 1.1, November 19, 1997 (OMG standard)

SIP–Based Call Centers: A Vendor-Independent Architecture for Multimedia Contact Centers

Neill Wilkinson
Senior Consultant
Quortex Consultants Ltd.

The current call-center market is populated with vendor-specific solutions, making the decision for the telecom manager on functionality very difficult. Finding the best fit and being able to combine components from different vendors can quickly become a major systems-integration project.

This paper outlines the possibility for a vendor-independent architecture that will allow best-of-breed components to be combined and allow for a close integration of systems without performing a major integration project. The architecture proposed also has the ability for rapid development of new applications for call centers by skilled Web developers (not systems integrators or expensive software houses) and provides a framework for eCRM products.

1. Introduction

During the last decade, computer telephony integration (CTI) has rapidly become the way in which call centers have achieved the benefits of being able to combine telephony services with information technology (IT) systems.

While this growth in CTI has improved customer care and call-center sales capabilities, the industry has been marred by proprietary systems and nonstandard interfaces, creating spiraling system-integration costs.

In the background to this has been the rise of the Internet revolution and the emergence of Internet protocol (IP) as the way to connect systems.

The growth of the World Wide Web (WWW) from Tim Berners Lee's original idea of a hyperlink information store to that of a powerful tool for application development has allowed for dynamic content to be generated and constructed based on inbound context.

Now that Internet telephony is becoming mainstream, the major telecommunications operators are looking to replace their circuit-switched infrastructures with packet-telephony networks [1].

This is allowing CTI, the WWW, and Internet telephony to be combined to create a new breed of call-center system, presenting an opportunity to utilize the open nature of Internet telephony for call centers and to move away from proprietary protocols and vendor implementations.

The call-center market has seen significant growth during the last decade. Recognition of the need to handle customers efficiently has lead to the ideas of customer-relationship management (CRM).

While the idea of customer relationship is not that new, the linking of customer relationship to the realization that retaining customers is more cost-effective than trying to win new ones is one that has captured the call-center industry.

The recent growth in the so-called e-economy has created a new area for customer contact, that of the Web and e-mail. While this has created a new channel to market, it has also created a new need to ensure good customer care through this medium. This has led to the growth in the eCRM market.

2. Evolution of CTI Architectures

The de facto architecture for computer telephony (CT) third-party call control is shown in *Figure 1*.

This architecture involves a CT server communicating with a private branch exchange (PBX) or automatic call distributor (ACD) system via a proprietary (vendor-specific) interface, and a proprietary interface between a workstation application and the CTI server.

The back office enterprise database is accessed through open database connectivity (ODBC), Java database connectivity (JDBC), or a direct structured query language (SQL)

FIGURE 1

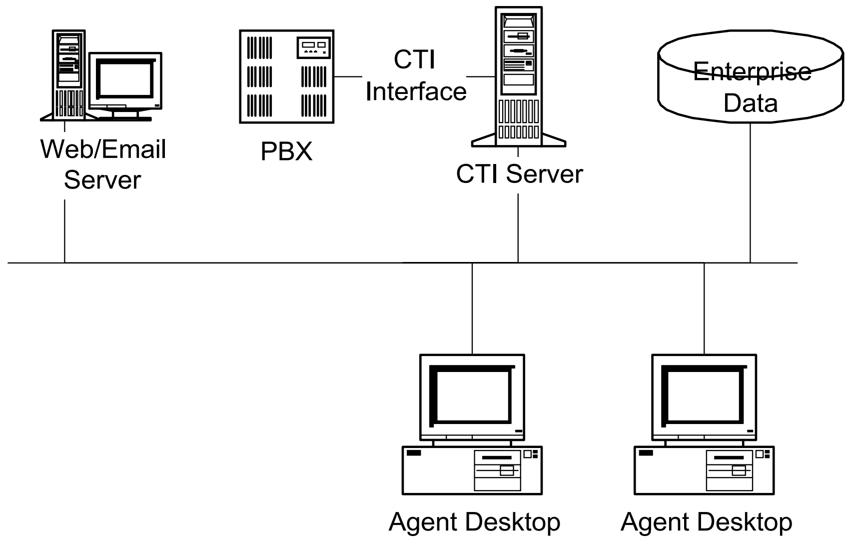

interface, either by workstation applications or by the CT server. This enterprise data is fed from marketing and other channels, and areas such as on-line analytical processing (OLAP) and data mining are very important in extracting value from this data. However, this is outside the scope of this paper.

This ends in a complex integration project of all the disparate systems through a plethora of interfaces and middleware, some of which is vendor-specific. This is clearly not an environment for rapid application development.

The Rise of the UnPBX

A little over three years ago, the use of NT as an operating system (OS) for call-center applications emerged as an opportunity to move away from proprietary PBX/ACD systems. "*We are total believers in the concept of Windows NT based UnPBXs. We fully expect that over time, smaller phone systems (PBXs and Key Systems) will be replaced by UnPBXs.*"[1]

A number of vendors seized this opportunity[2] and started to produce hardware and software alternatives to the box-based ACDs/PBXs. The UnPBX was born [3].

The rise of the UnPBX was the birth of the movement away from circuit-switched proprietary ACDs to a more open IP–based solution. However, three years ago (and to an extent now) proprietary call-control protocols still reigned supreme[3].

The current trends in open ACDs owe a lot to these UnPBX vendors. The telecom world has moved on from there, and a more open carrier-sized upgrade is taking place.

3. Next-Generation Networks

A lot of words have been printed about voice over IP (VoIP) and IP this and IP that. The plain fact is that IP has become the de facto standard for multimedia communications. This means that ways have been sought to carry voice *reliably* over IP. This has lead to the emergence of what are known as the next-generation networks. What this means for call-center technology is that there is a greater opportunity to capitalize on the changes in the core network.

The New PSTN

Conventional circuit-switched time division multiplexing (TDM) networks are slowly being replaced by packet-switched networks carrying both voice and data. These networks are being created from new access techniques, such as digital subscriber line (xDSL) and hybrid fiber/coax (HFC) cable modems, and core carrier network technologies such as media gateway controller (MEGACO) [4]. The mobile networks have not been left out, and the Universal Mobile Telecommunications System (UMTS) standards are based on an asynchronous transfer mode (ATM) transport, with increased focus being placed on data services (general packet radio service [GPRS] and enhanced data rates for GSM evolution [EDGE]; and wireless application protocol [WAP] and i-Mode).

While this creates the opportunity for more integrated applications, it will also force the telecom operators to consider redesigning all of their conventional telecommunications services.

Since most (if not all) major telecom carriers operate separate voice and data infrastructures, major consolidation will be required to efficiently move the voice network from TDM to the packet infrastructure of the data networks. To get real value (read revenue) from this infrastructure investment, without pricing the services out of everyone's reach, the telecom operators will need to look at technologies such as virtual private networks (VPNs) for creating virtual segregation of customer traffic.

There is a major market opportunity for a new breed of telecom operator to emerge, the *IP telco*. This operator could emerge from an existing network operator or could be launched from a global Internet service provider (ISP) like AOL. The IP telco could operate in a model similar to that of the mobile virtual network operator (MVNO) or switchless resale company.

To capitalize on the VNO model, incumbent and competitive global telecom operators could deliver telecommunications services on a wholesale basis from their networks, rebranded by VNOs completing the evolution of the switchless resale model.

The cost of entry for a new player to come into the market place as a global IP telco is relatively small, and the potential for large revenues is huge[4]. The revenue is rapidly becoming derived from the services offered, not the cost of access to the services[5].

This means a new entrant could enter the marketplace without having to own any network, they *just* have to negotiate peering access to a number of key points of presence (POPs) around the world and piggy-back their service on other incumbent networks. Clearly, new peering arrangements that include quality of service (QoS) and service-level specifications will have to replace the current ISP peering model.

Voice-Based Information Services and the New Answering Machine?

A number of the main vendors of interactive voice response (IVR) systems are looking to add packet voice interfaces to their products.

This is being augmented with extensible markup language (XML)–based scripting in the form of VoiceXML[6]. This allows voice prompt scripts to be written in a platform independent language that also allows a close integration with Web servers.

Unified communications (UC) is an application that has received a great deal of press recently. UC is unified messaging (UM) come of age. It provides the ability to integrate voice-mail, e-mail, fax, short message service (SMS), and paging into a single inbox. Work has been done on the use of SIP and real-time streaming protocol (RTSP) for a UM platform [5].

UC looks set to replace company voice-mail systems and to provide a new answering-machine service to users of in-net services via an application service provider (ASP) delivery model. British Telecom (BT) offered an early version of an in-net answering machine with their CallMinder™ product.

A number of ISPs are starting to offer UM services for combining fax and e-mail integration and allowing customers to dial in and listen to their e-mail. The problem with these at the moment is that they require their users to have another access number.

4. The Vendor-Independent CTI Architecture

Building on the foundations of the UnPBX, the softACD has appeared, and companies such as CosmoCom[7] are successfully utilizing software frameworks—such as Microsoft distributed component object model (DCOM), J2EE, and common object request broker architecture (CORBA)—to develop conventional ACD functionality and combining Web server capability to deliver eCRM.

These first-generation softACDs are creating a new opportunity for vendor-independent CTI capability by combining the CT server and ACD into a single platform built from open software frameworks.

The architecture of a truly vendor-independent solution is available now through products such as Sun Microsystems' J2EE, Oracles' 8i or IBM's Websphere, and BEA's Weblogic.

What vendors will do is to write code on top of these application frameworks, providing the value-add for call-center applications. Application frameworks such as this are also big business, and companies such as BEA Systems have made significant inroads into this market segment.

The Open-Source Movement

By combining open-source H.323 and/or session initiation protocol (SIP) stacks, anyone with software engineering skill can develop a softACD and provide Web integration to deliver an eCRM solution. The biggest challenge is adding the value through reporting tools and remote maintenance capability and combining these with sound back-office integration.

This is creating an opportunity for Web developers to create new services, building on top of platform libraries without having to resort to help from systems integrators. This also means that small businesses need not worry about large bills from system integrators and can take services from network-based vendors for a fraction of the cost.

This also opens up the possibility for large telecom operators to take the code developed by the new entrants to the market and transport it the large network-based call servers, thus delivering large-scale, feature-rich, network-based softACDs. This also creates the opportunity for ASP delivery of customer contact systems for the new e-economy.

It also opens up the way for systems/software houses to enter third-party development and partnership agreements with large operators to allow for the rapid delivery of new differentiated services.

While the traditional role of systems integrators is diminishing, their role is evolving from technology integration to business integration—moving them up the value chain.

All of this software is of course nothing without the call-control protocol. SIP [6] is emerging as the favorite to fill this space. However, H.323 is probably currently the most prevalent implemented call-control protocol. For comprehensive comparisons of these protocols refer to [DALGIC].

Sun Microsystems is working on a Java-based model called Java APIs for integrated networks (JAIN) that incorporates these architectural ideas, including a SIP and H.323 stack for intelligent network (IN)–based services.

The protocols responsible for converged communications are best viewed as a layered model [7]. This model consists of the call-processing protocols, user protocols, and support protocols.

Call-Control Protocols
The call-control protocols have been highlighted as SIP, H.323, and MEGACO.

Of these, SIP is arguably the most versatile and extensible, as it is text-based and designed around simplicity. This, however, does leave out a lot of embedded feature capability.

Why Use H.323?
H.323 is actually a collection of protocols, and H.245 is actually the call-control protocol (based on Q.931) combined with H.225. This framework has evolved from the work on the H.320 videoconferencing standards, which means a lot of capabilities and features are supported by the H.323 family of protocols.

Also, H.323 is important because of its prevalent implementation by Microsoft in its Netmeeting' product and in the TAPI 3.0 implementation on Windows 2000.

SIP's Attractiveness for Call Control
One of the key issues of any Internet-based service is security. SIP can support security through the use of any encryption standard. For example, the widely available pretty good privacy (PGP) package could be used to encrypt the signaling payload. One of the potential issues with SIP is that IP addresses are carried in the header message; this can be overcome with a SIP–network address translation (NAT) module (a version is available for Linux under the GNU license). Solutions to this are far from elegant—yet [8], and it may be necessary in a network operator's environment to implement a firewall control protocol.

The SIP architecture has two key components that provide for load balancing and service transparency, namely proxy servers and redirect servers for load balancing.

SIPs openness and extensibility allows for easy extension. The SIP INFO method [9] for in-call messaging is an example of this and is a perfect candidate for in-call communication between an agent desktop and a proxy server softACD.

Support Protocols
The list of support protocols is endless, as this is essentially the list of all the specifications of the Internet Engineering Task Force (IETF). That said, the protocols of most significance are IP multicast for agent monitoring and third-party call connect; the IP security (IPSec) specifications for the creation of IP–VPNs and multiprotocol label switching (MPLS) for the same purpose; the RTSP and real-time conferencing protocol (RTCP) for playback of recorded messages; and finally the Internet relay chat (IRC) client protocol for real-time chat between customers and agents.

One could bundle in the plethora of Web protocols, such as hypertext transfer protocol (HTTP); however, these are more related to the Web-server delivery of content, rather than collaborative interaction.

Application Frameworks
A word of caution on the title of this section: Application frameworks are an area of study relating to the use of object-orientated techniques for the production of domain-specific software components. I do not intend to expand on this here, as this is a massive topic area in its own right. If application frameworks interest you, I suggest you refer to [MOHAM].

This section highlights some of the middleware frameworks that are available from the development of the collaborative applications necessary to support an eCRM service.

Firstly, Sun Microsystems' J2EE coupled with its iPlanet Web server. These products build on Sun's Java 2 virtual machine for the rapid development and deployment of component-based applications and create an environment for the creation of an integrated ACD call server.

The BEA/Nokia and Motorola (MIX) platforms use the same capabilities as Java server pages and Java servlets to create an integrated environment for wireless platform application development. These could be used as the basis of a wireless communications server.

Finally, dynamicsoft's proxy server products utilize CPL and SIP CGI for call control in a SIP environment. It is a small mental step to enhance this to a Java servlet environment to provide softACD functionality on top of this platform and to link these platforms to Web servers to provide multimedia integration. Queue and agent skill scripting could be provided by CPL [10,11].

Call Servers
The evolution of the public switched telephone network (PSTN) from a collection of circuit switches to a distributed call server (media gateway controller) and gateway (media gateways and signaling gateways) architecture, and the use of proxy servers in the SIP architecture, has created the opportunity to deliver all services from call servers.

This might be considered as an issue with "all your eggs in one basket," but that is not the case, because the call servers are in their own right fault-tolerant, distributed platforms.

What is key is that the call servers are being constructed from *standard* IT hardware. This opens up the possibility for these servers to be used for network-based softACDs.

Billing for Network Services
Clearly, from a network operator's perspective, being able to charge for the use of a service is very important. Billing for next-generation network services will be performed by the call servers (media gateway controllers, softswitches, SIP proxy servers, H.323 gatekeepers). The difficult arises from the fact that once an endpoint has discovered the other party's endpoint, there is theoretically nothing stopping the originating endpoint from passing the call server, thus creating the opportunity for fraudulent use of network resources.

To prevent this in, for example, an SIP situation (a similar solution would work for H.323 endpoints), the SIP user agent must contact the SIP proxy via a NAT firewall, thus preventing any visibility of the "internal" network [12]. This of course leads back to the issues involving the use of firewalls.

5. Conclusions

Firstly, let's start with areas not covered by this paper: business models (marketing/sales/service sale cycle, when things go wrong, order tracking, workflow) and delivery models (ASP, customer-premises equipment [CPE]). These areas are papers in themselves, and the author suggests that the reader look elsewhere for discussions of these topics.

Telecom operators need to be looking at platform vendors rather than mega-software vendors.

The mega-software vendors will develop functionality based on what the mass market wants, and customization and differentiation will be more difficult than with a platform solution.

The platform solution has the advantage of allowing for vertical market products on a common architecture. The scale of product offering required by a telecom operator makes the platform solution the obvious choice. For wholesale delivery of services to different customers, telecom operators will also need to consider the ability to segment the services in a secure fashion.

Watch out for consolidation in this market space. The market is still young, and vendors such as Cisco Systems, Alcatel (Genesys products), Siebel, and Clarify (a Nortel Networks company) have the power to swallow new entrants. The author has specifically steered clear of the CRM products provided by companies such as Siebel and Clarify, purely because the space that these vendors occupy is currently the area of product and system integration. This white paper's whole premise was to discuss the framework that will underlie these vendors' products.

Also watch out for the platform people (BEA, Vignette, and Art Technology Group) and for the architecture vendors (IBM [Websphere], Sun Microsystems [J2EE and iPlanet]. The other big boys keen to push this opportunity are, of course, Oracle and Sybase. Oracle's 8i and 9i products could form the core of many eCRM offerings.

Another company in this space that will later emerge as an eCRM candidate is dynamicsoft. dynamicsoft is the leader in SIP–based products that will underpin the voice side of softACDs. (The company's chief scientist, Jonathan Rosenberg, is very active in this area and has recently proposed application architecture that may go someway toward achieving this [13].) All of the other vendors mentioned are firmly from the IT and data-management camp. So far these IT vendors have not integrated the voice portion of customer interaction, arguably the most important channel.

6. References

1. "Telecoms Operator Aims for Switch to Internet Technology within 3 Years," *Financial Times*, October 3, 2000.
2. *eCRM: Stop the Insanity*, UBS Warberg Global Equity Research.
3. Edwin Marguiles, "The UnPBX," (Flatiron Publishing, Inc.), ISBN: 1-57820-014-8.
4. RFC2885, MeGaCo Protocol.
5. Kundan Singh and Henning Schulzinne, *Unified Messaging Using SIP and RTSP*, Columbia University.
6. Rosenberg et al., latest draft of RFC2543, expires January 2001,draft-ietf-sip2543bis-02.
7. Uyless Black, *Internet Telephony (Call Processing Protocols)*, (Prentice-Hall), ISBN: 0-13-0225565-3.
8. J. Rosenberg and H Schulzrinne, "SIP Traversal through Residential and Enterprise NATs and Firewalls," draft-rosenberg-sip-enfw-00.txt.
9. Steve Donovan, RFC 2976 SIP INFO method, dynamicsoft.
10. Jonathan Rosenberg, Jonathan Lennox, and Henning Schulzinne, "Programming Internet Telephony Services," Tech Report Number CUCS-010-99.
11. J. Lennox, J. Rosenberg, H. Schulzrinne, draft-lennox-sip-cgi-01.txt, Columbia University/Bell Labs.
12. Olivier Hersent et al., *IP Telephony Packet-Based Multimedia Communications Systems*, Addison Wesley.
13. Rosenberg, Mataga, and Schulzrinne, "An Application Server Component Architecture for SIP," draft-rosenberg-sip-app-components-00.txt.

[DALGIC]
Ismail Dalgic, Hanlin Fang, "Comparison of H.323 and SIP for IP Telephony Signaling," 3Com Corporation, Technology Development Center, 5400 Bayfront Plaza, M/S 3219, Santa Clara, CA 95052; Cisco Systems, 170 West Tasman Dr., San Jose, CA 95134.

[MOHAM]
Mohamed E. Fayad, Douglas C. Schmidt, and Ralph E. Johnson, *Building Application Frameworks – Object Oriented Foundations of Framework Design*, (Jon Wiley & Sons), ISBN 0-471-24875-4.

Notes

1. Bill Gates, Microsoft, *Computer Telephony Magazine*, June 1997.
2. Dialogic, Brooktrout, Netspeak, et al.
3. Nortel – Compucall & Meridian Link, Aspect Application Bridge and Event Bridge, Lucent ASAI, Netspeak Communications Protocol (NSCP), et al.
4. In economic theory, in a perfectly competitive environment, if excess profits exist and barriers to entry are low, competition will eventually appear and erode excess profits.
5. Local call charges are rapidly being dropped for Internet access, and broadband is being deployed on fixed fee packages.
6. a.k.a. VoxML
7. http://www.cosmocom.com

Voice over Internet Protocol Quality of Service

Bruce Young
Engineering Vice President
Telverse Communications, Inc.

Philip Repino
Publications Manager
Telverse Communications, Inc.

An opportunity to use data networks for voice over Internet protocol (VoIP) telephone traffic holds great appeal because VoIP can reduce the business expense of long-distance calls and offer enhanced features. Today, VoIP is an evolving technology that is gaining in popularity and application. VoIP is evolving toward real-time conferences, presentations, and video streaming. However, as more companies deploy VoIP, more Internet protocol (IP) networks will have to carry time- and jitter-sensitive packets. A circuit-switched public switched telephone network (PSTN) and a packet-switched Web network hold different challenges. IP networks require testing and monitoring for call set-up, voice quality, and packet analysis. Additionally, while bandwidth consumption may be relatively low, VoIP poses stringent demands for low latency and consistent arrival of voice-related datagrams. Each voice carrier and enterprise should calibrate its network for optimal VoIP performance.

VoIP Simplified

As stated, VoIP eliminates long-distance telephone charges by sending voice signals over the Internet. Previously, the use of the PSTN incurred time-based, long-distance charges. By using VoIP, packets of digitized voice can now be sent over the Internet to a distant point within the VoIP Internet, or a call can be switched upon receipt to the local PSTN and onward to any digital or analog telephone.

Data can travel over a packet-switched network in packets that arrive at varying times. They are simply reassembled back into their correct sequence when received. However, for real-time voice transmission, hearing packets of digitized voice in anything but the correct order is unintelligible. Re-assembly techniques do not work for voice packets.

Quality of service (QoS) can be improved by using software to compensate for many of the problems of sending voice over packet-switched networks, such as delay, jitter, and random or recurring lost packets. With this new software and other tuning adjustments, packets of digitized voice can be sent successfully over a data network.

The Inherent Parts of VoIP: Delay, Jitter, Echo, and Packet Loss

Delay and Jitter

Because network congestion can be encountered at any time within a network, buffers can fill instantly and cause packet delays. Jitter is the difference between when a packet is expected to arrive and when it actually is received: the end-to-end delay. (Jitter refers to *packet* jitter, not bit jitter). To compensate for delay variations between voice packets in a conversation, VoIP endpoints use jitter buffers to store incoming packets and release them in a more constant stream and thus turn the delay variations into a (unperceivable) constant value. This damping allows a voice to be delivered to a handset smoothly. A suggested goal is to keep jitter at less than 20 ms, ideally at 5 to 10 ms.

As jitter buffers hold voice packets, they add latency. And, compounding the jitter problems somewhat, packets can get lost if a jitter buffer overflows. Thus, the size of a jitter buffer can affect both jitter and latency. If voice traffic has enough jitter to annoy users, increasing the size of a jitter buffer will reduce jitter to acceptable levels. Too large a buffer, however, may cause latency to increase to where it is again annoying to users.

There is no optimal size for a jitter buffer, and buffer size varies within the network. A typical jitter buffer delay is 20 ms but can reach 80 ms. Increased bandwidth and proper tuning can reduce network congestion, which in turn will reduce packet loss and jitter.

Several considerations and possible actions exist to counteract delays. Many types of delay can present themselves in today's network, such as jitter buffer (as mentioned), network-appliance transport, propagation, and packetization. To compensate for delay added by appliances, the number of hops and the delays within the worst-offending appliances can be reduced. Propagation pitfalls can be reduced by making the network path as direct as possible. Packetization delay can be reduced by using a coder-decoder (CODEC) compression scheme that induces only a small packetization delay. Finally, not allowing changes from one CODEC to another along a path and avoiding complex, time-consuming CODECs will reduce delay.

Latency

Designs for VoIP transmission have to measure the amount of latency that an appliance will add to a network. By knowing the latency of each component, you will know what each component contributes most to the total latency. That information can help you to decide if the size of buffers is sufficient to deliver acceptable QoS.

If network latency is too great, VoIP conversations degrade. Typically, users can tolerate delays of no more than about 200 to 250 ms before a conversation becomes annoying. Delays of 400 ms to 500 ms make conversations impractical. The suggested goal is to keep latency less than 150 ms for one-way traffic between endpoints.

Echo

In a voice telephone call, echo occurs when you hear your own voice repeated. An echo is the audible leak-through of your own voice into your own receive (return) path. Echoes arriving before 20 ms are generally imperceptible—those arriving after 20 ms are audible.

Some things that can be examined to reduce echo include the following:

- Check speakerphones or headsets. If the destination telephone is a speakerphone or headset, it is probably the source of the echo. Replace the speakerphone or headset with a new or better-quality handset and see if the echo vanishes.
- Telnet to the destination voice gateway and check that the echo canceller is enabled and that the coverage is set to maximum.
- Test for normal echo canceller. If the echo persists and you have verified that the echo canceller is working properly, the echo canceller cannot fix the echo because it is too loud (called a *loud echo*) or if it is too delayed (called a *long echo*). Most persistent echoes are loud echoes. Delayed echoes are common, however, when the end circuit involves a long-distance PSTN link, a series of alternating digital and analog links, or any other link with high latency.
- Match output impedances and levels with the analog telecom equipment attached to the adapters and gateway analog voice ports.

Today, echo cancellers, built to the International Telecommunication Union (ITU) G.168-2000 standard, can provide rapid (< 500 ms) training and some degree of real-time echo cancellation for telephony applications.

Packet Loss

Heavy traffic loads are a primary contributor to network delay, jitter, and dropped packets. When a router overloads with traffic, it may intentionally drop packets to relieve congestion. A voice packet that arrives late might as well have not been transmitted at all. Retransmission is of no value. And, excessive packet loss affects the QoS of a VoIP system. Because lost packets cannot be recovered, they show up as gaps in a conversation. Small gaps are acceptable, but a consistent high rate of loss is not acceptable.

Typically, more than 5 percent of lost packets will annoy users. Burst packet loss is more obvious to users than random packet loss for the same rate of loss. If the average packet loss rate is 1 percent, but the dropped packets are widely dispersed, quality can be perceived as quite good. Gateways that provide packet-loss monitoring but do not consider the difference between burst and random loss will give misleading results.

VoIP using G.711 CODEC delivers 55 packets per second. With a 5 percent loss of packets that results in a reception of about 52 packets per second (pps), an acceptable value still exists. Protocol analyzers will display the number and percentage of lost packets, and analyzers can tell you when lost packets occur in bursts. For example, if you lose every 20th packet (5 percent), the conversation is still understandable. The goal is to keep random packet loss well below 1 percent, ideally at approximately 0.2 percent.

One clever practice to reduce the effect of random packet loss is called *packet-loss concealment* (PLC). This masking technique interpolates a portion of the information from the last good packet and substitutes it for the missing packet, making the loss unnoticeable to the user.

Firewalls

Getting session initiation protocol (SIP) and media gateway control protocol (MGCP) through firewalls and network address translations (NATs) is particularly troublesome. Generally, firewalls are configured to not let user datagram protocol (UDP) packets in or out. (The fields related to real-time transport protocol [RTP] are inside the UDP.) Transmission control protocol (TCP) packets are not allowed in unless they are destined for specific servers designated to handle the protocol. Some UDP–entry solutions are based on having a proxy server control the firewall/NAT with a control protocol. This protocol can open and close holes in the firewall and obtain NAT address bindings to use in rewriting the session description protocol (SDP) in a SIP message.

Firewalls are beginning to appear that operate on a protocol-inspecting basis and allow IP telephone signaling control of the firewall, IP telephone file transfer protocol (FTP) configuration, and NAT look-ups. In yet another approach, application-layer gateways can serve as external voice gateways to forward voice on a separate access network. The ALG functions as a proxy and acts as a firewall. The ALG must assume a limited set of router capabilities to support registered network appliances. This simplifies QoS and legacy firewall issues that present themselves in a shared data and voice deployment.

RTP Header Compression

RTP is widely used to encapsulate video and audio frames because it is designed to send audio quickly in one direction without acknowledgement. Each RTP header contains a timestamp that enables the receiving end to reconstruct the timing of the original voice data and deal with duplicate, missing, or out-of-order packets. RTP header compression, known as cRTP, can reduce a 40-byte RTP header by 90 percent and thus reduce bandwidth use. As with jitter buffering, cRTP adds latency and is best used on links with 500 kilobits per second (kbps) bandwidth or less.

Differentiated Service for VoIP

Because VoIP has strict requirements for packet loss, delays, and jitter, it makes sense that it be treated with a high QoS. One method is to add the differentiated services (DiffServ) tag to VoIP packets for expedited flow (EF). DiffServ is a QoS tuning technique that assigns a high priority to RTP packet streams and is intended to improve handling as the packets pass through routers, switches, and other network appliances. This technique, however, does not always deliver guaranteed results. Handling methods and policies can vary among network devices, and what began as a voice stream bearing a higher setting than best effort can sometimes be downgraded or re-queued with other high-priority traffic. Remember, too, that inserting packet tags introduces delay with an accompanying jitter.

Other QoS Issues

The greater part of this paper has focused on the factors that can affect voice transmission using VoIP once a call has been placed. There are other endpoint QoS issues that need to be considered. Call–set-up time, set-up rates, dropped calls, and success ratios while using VoIP present performance challenges and need to be measured and maintained.

Voice-Quality Standards

Standard groups, such as the ITU, address the issue of voice quality and have published two important recommendations: P.800 (mean opinion score [MOS]) and P.861 (perceptual speech quality measure [PSQM]). P.800 defines a method to derive a voice-quality MOS, which is scored from 1 (bad) to 5 (best). A MOS of 4 or higher is considered carrier-grade. P.861 defines an algorithm that a computer can use to derive a score that has a close correlation to a MOS score. PSQM does not consider factors such as jitter or frame loss. Keep in mind, though, that intrusive snapshots that gather voice-transmission averages can actually impair quality by using valuable bandwidth themselves.

Another standard has emerged for VoIP from British Telecom called the perceptual analysis measurement system (PAMS), which produces a good correlation between itself and manual MOS results. The European Telecommunications Standards Institute (ETSI) is another organization that produces speech-quality standards.

In the final analysis, the quality of VoIP hinges on a positive perception of subjective people using telephones.

Summary

Most likely, your data network will have to be tuned to carry VoIP traffic. The majority of data networks contain points of congestion and jitter, but with today's network-testing tools, it is easy to test a network's readiness or locate bottlenecks by merely simulating VoIP traffic. This allows you to make needed network changes to assure VoIP success. Examine the network frequently under varying conditions until you are convinced it is ready for voice. Later, re-examine performance as you add services.

Section V:

Wireless

Safety Standards for Wireless Optical Networks

Jim Alwan
Manager, Electro-Optics Group
AirFiber, Inc.

Introduction

Wireless optical networks (WONs) are beginning to emerge in the telecommunications market as a strategy to meet last-mile demand, enabling reliable high-bandwidth connectivity previously available only to customers directly connected to fiber. Open-air or free-space beams used by WONs to transmit data are not a new concept. Historically, such free-space systems were used in point-to-point (PTP) connections in some campus settings and by the military and the aerospace industry. Expanding into the commercial sector, this technology has evolved into systems with back-up and redundant optical links, providing high reliability and fiber-like bandwidth to customers located up to a kilometer away from buried fiber. Such systems are being deployed to commercial buildings in urban areas, breaking the so-called "last-mile" bottleneck. These WONs provide a unique solution that provides higher bandwidth than radio frequency (RF) wireless systems and is considerably less expensive than laying additional fiber. Additionally, there are no complex licensing requirements when implementing WONs.

The use of wireless optical networking technology enables carriers to offer low-cost, high-bandwidth services that can be deployed within a short period of time. However, implementing these systems in public spaces provides new safety challenges. Understanding existing eye safety standards, as well as the factors in the WON environment that impact eye safety, is essential in addressing these challenges.

Equipment makers, laser-safety professionals, and telecommunication service providers are working together to build safe and reliable systems that deliver the next generation of broadband connectivity. This whitepaper presents a review of laser ocular hazards, present and upcoming safety standards, and eye-safe WON architectures. Additionally, the paper highlights considerations when evaluating the eye-safety of wireless optical networks and uses the AirFiber OptiMesh Network as an example of a Class 1–compliant WON.

Light, Lasers, and the Human Eye

The human eye is a sophisticated and sensitive optical detector with a peak visual response to light in the visual range (400–700 nm)[1] that has evolved over time to correspond to the peak of the solar spectrum. Like any other living tissue, the tissues of the eye are susceptible to damage under extreme conditions. Because different tissues absorb some wavelengths more strongly than others, different parts of the eye are variably susceptible to damage at a given wavelength. The cornea (glassy front of the eye), the lens (behind the cornea and in front of the retina), and the retina (the back surface of the eyeball that detects light) are all in the path of light entering the eye and therefore may be subjected to dangerous light levels.

Many factors combine to result in the potential for laser-induced ocular injury. The anatomy of the eye, physiology of vision, and laser-tissue interaction are important factors. Laser-tissue interaction is strongly dependent on wavelength, power, and pulse duration. Damage can occur by several mechanisms including photochemical, thermal, and thermo-acoustic. *Photochemical* interactions are those in which an absorption of a photon by a molecule results in a chemical reaction. *Thermal* interactions refer to the deposition of heat to a local area, whereas *thermo-acoustic* refers to transient phenomena where the rapid deposition of heat results in a damaging shock-wave. In terms relative to the human eye, a thermal interaction might result in a slow burn at the irradiated site, whereas the thermo-acoustic interaction could cause additional physical damage beyond the burn site, such as detachment of the retina.

As described, specific tissues in the eye interact strongly with different portions of the optical spectrum. The cornea and lens are transparent to the visible wavelengths and near-infrared, as they are intended to focus visible light onto the retina. However, ultraviolet (180–400 nm) and mid-infrared (1400 nm–3 µm)[2] to far-infrared (3 µm–1 mm)[3] exposure to the cornea and lens can cause photokeratitus and cataracts, respectively. *Photokeratitus* (also known as welder's flash or snow blindness) is similar to sunburn on the cornea. This condition, while painful, does not normally result in permanent damage. *Cataracts* refer to the formation of cloudy regions in the otherwise clear lens, compromising vision as they prevent light from reaching the retina. Visible and near-infrared light is of particular hazard to the retina as this is the range of light to which the eye is intended to be sensitive and thus has a high absorption. Hazardous exposure to radiation in this range can result in retinal burn resulting in compromised vision at the damaged site.

Lasers are potentially hazardous as a result of their high brightness (also known as radiance, W/m^2 sr). A handheld laser pointer (~630–650 nm and visibly red) is many times brighter than intense lamps or even direct sunlight. Although a lamp or the sun may have a larger total output of light, it is generally spread into a large solid angle. For example, the sun shines in all directions around it and is thus said to be radiating into 4 steradians (sr) of solid angle. In direct contrast, a laser can produce a beam of light that spreads very little and radiates into a solid angle several thousand times smaller; thus the brightness is several thousand times higher. It is the collimated, beam-like quality of laser output that results in very high irradiance (also known as power density, W/m^2), because lasers can be focused to a much smaller spot on the retina than conventional light sources. *Figure 1* provides a comparison of the retinal irradiance of a variety of sources.

Laser Safety Standards

The eye-safety hazards of lasers were recognized within a year of the first demonstration of a laser device. As some of the first adopters of laser technology and research, military agencies issued the earliest laser-safety guidelines. In the 1960s and '70s, biomedical research regarding laser-tissue interactions established the first physiological basis for safe exposure limits. In more than 40 years of research and development, the variety of lasers now spans the spectrum from the vacuum ultraviolet to the far-infrared, with pulsewidths as short as femtoseconds (10^{-15} s). The task of providing laser-safety guidelines and regulatory codes that adequately address such a broad field of devices and diverse applications has become significant. Just as the capabilities of laser systems have expanded, so has the complexity of safety standards derived to address them.

There are a number of regulatory, governmental, and educational organizations that have developed and published laser-safety standards, three of which are discussed in the following paragraphs. The standards generally have the following main functions:

- Classification of laser systems
- Safety measures for manufacturers
- Guidelines for safe operation by users

Some standards are for manufacturers of laser systems providing regulatory codes on how laser products can be made, classified, and labeled, for example. Other standards are considered user standards and address how individuals and organizations should deploy and safely use laser systems in the community and the workplace.

ANSI Z136.1–1997 The Safe Use of Lasers

One of the earliest standards that has undergone continual revision since its inception in 1973 is American National Standards Institute (ANSI) Z136.1–1997 *The Safe Use of Lasers*. This document provides classification based on maximum permissible exposure (MPE) in detailed chart, graph, and equation form, expressing such limits as a function of wavelength and pulsewidth. It is not regulatory, but rather a user standard. It has been adopted by corporations,

FIGURE 1

Retinal Irradiance for a Variety of Sources
The retinal image area influences the hazard as does the source radiance.

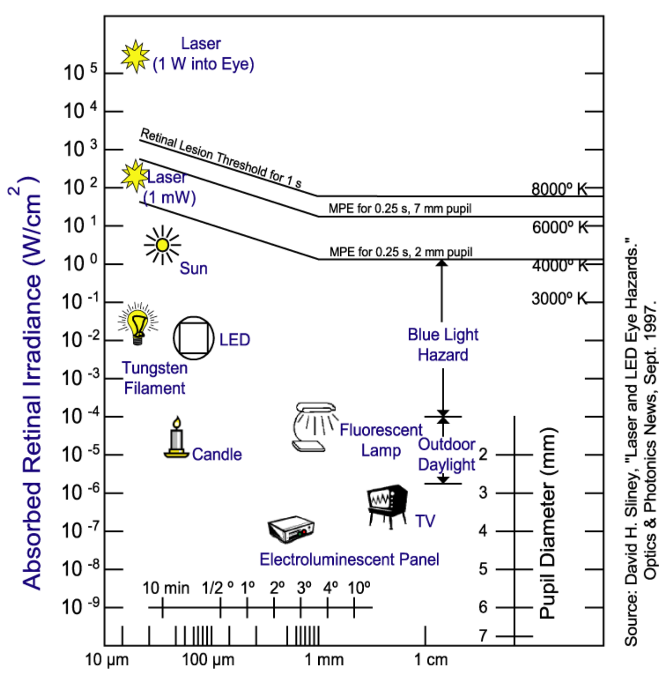

municipalities, and some states in the United States as an official standard for safe use and deployment of laser systems. It is aimed at health and safety professionals charged with providing a clear and well-grounded basis for the use of laser systems as they become more prevalent in our everyday life.

IEC 60825-1 (2000) Safety of Laser Products Part 1–Equipment Classification, Requirements, and User Guide

The International Electrotechnical Commission (IEC) IEC 60825-1 (2000) *Safety of Laser Products Part 1–Equipment Classification, Requirements, and User Guide* is an internationally drafted and recognized technical standard for users and manufacturers. This standard serves as the basis for a subset of standards targeted toward more specific applications—for example, the IEC 60825-2 *Safety of Optical Fibre Communication Systems*. The IEC standards are often directly adopted by countries around the world—or at least form the basis for national policy and regulation on the subject. For manufacturers of products that are sold internationally, compliance with the IEC standards is requisite. In contrast, the United States has an independent standard under the Code of Federal Regulations (CFR).

21 CFR Ch. 1 (4-1-97 Edition) Part 1040 Performance Standards for Light-Emitting Products

The U.S. Food and Drug Administration (FDA) includes the Center for Device and Radiological Health (CDRH), which administrates the U.S. regulatory laser standard. That standard is the 21 CFR Ch. 1 (4-1-97 Edition) Part 1040 *Performance Standards for Light-Emitting Products* Sections 1040.10 and 1040.11. This standard stipulates a classification system, manufacturing and labeling guidelines, and submission of a report of compliance with the CDRH.

Table 1 provides a summary of laser-safety standards currently in existence.

Classification Systems and Certification

Classification

There are three concepts that are normally included (directly or by reference) in a laser-safety classification system: class definitions, accessible emission limit (AEL), and MPE. The class definitions provide non-technical descriptions understandable to the lay-person; the AELs define the classification breakpoints; and the MPEs are based on biophysical data and indicate actual tissue-damage thresholds.

The classification allows an abbreviated way to readily communicate the hazard level to a user by means of classes 1, 2, 3, and 4, for example. It should be noted that various standards, including those described in the previous section, use confusingly similar yet distinct classifications in the form of Arabic and Roman numerals and upper- and lower-case English letters (e.g., Class 2, II, 2A, 3b). These may appear on a user manual or label, for example. The classes are described in words to indicate the general hazard posed by a laser in a given class. *Table 2* shows the classification system of the IEC 60825-1, with the word description and AELs for an example laser of a doubled CW Nd:YAG laser (μ = 532 nm). This is a simple example. The addition of pulsed modulation, extended sources, multiple wavelengths, and so forth are treated with correction factors. Recently added classifications in the IEC60825-1(2000) include Classes 1M and 2M. These classes take into account that applications exist where aided viewing is not likely. In such cases, the AEL is considered for the unaided viewing case only, and therefore Class 1M (2M) is defined as a less stringent class than Class 1(2).

The AEL refers to the power level at a given wavelength that signifies a laser as belonging to a particular class. Generally, a standard includes tables, graphs, and equations of AELs that allow a person to classify a given laser based on wavelength, power, and pulse duration. The AEL may be related by some safety factor to the MPE values.

The MPE values are based on actual biophysical research of laser-tissue interactions spanning wavelength, power, and pulse duration for a specific tissue. They are defined as the level of laser radiation to which, under normal circumstances, persons may be exposed without suffering adverse effect. The MPE are related to the wavelength, pulse duration, and, in the case of visible and infrared, the size of the retinal image. There is a considerable body of research supporting the MPE. As lasers reach new benchmarks of performance, particularly in pulsewidth and power, research is performed to estimate the hazard presented. It should be noted that the MPE values are generally derated by a safety factor themselves, adding further margin to MPE–based AELs. In summary, it should be understood that class limits in the form of AELs are not tissue-damage thresholds, but rather are significantly reduced below such limits in the interest of safety.

Certification

A Class 1–certified laser is considered eye-safe with no special precautions required for persons who may be exposed

TABLE 1

Summary of Existing Laser Safety Standards

Standard	*Type*	*Target*
ANSI Z136.1	User	Laser users and health and safety professionals
IEC 60825-1	Technical	Manufacturers with global distribution
21 CFR Ch. 1 (4-1-97 Edition) Part 1040	Regulatory	Manufacturers in the United States

TABLE 2

IEC 60825-1 (1998) Classification of Doubled CW Nd: YAG (μ = 532 nm)

Class	*Description*	*AEL*
1	Lasers that are safe under reasonably foreseeable conditions.	0.39 μW
2	Visible lasers (400–700 nm) where aversion response such as the blink reflex affords eye protection	1 mW
3a	Lasers that are safe for viewing with the unaided eye. Direct viewing of the beam with aids (e.g., binoculars) may be hazardous.	5 mW
3b	Lasers that are hazardous when the beam is viewed directly.	500 mW
4	Lasers capable of producing hazardous diffuse reflections and that may pose skin and fire hazards.	> 500 mW

Note: microwatts (μW, 10^{-6}W) and milliwatts (mW, 10^{-3}W) are measurements of power used in the AEL

to the system. Being Class 1 requires that a system comply with all Class 1 requirements in the applicable standard(s). Most generally, this means that the system must limit access to radiation below the Class 1 AEL under any reasonably foreseeable single-fault condition. As previously indicated, AELs for a particular class definition are wavelength-dependent.

The CDRH CFR 1040 and the IEC 60825-1 are self-regulated standards in the sense that the manufacturer has the responsibility to correctly interpret and apply the standards to products. There is no mandatory U.S. federal or international compliance testing. The perils of not fully understanding and complying with applicable standards are manifol, and include product liability and, most importantly, potential hazard to humans.

There are a number of nationally and internationally recognized, independent laboratories that perform laser-product safety testing. Examples include Underwriters Laboratories (UL) and TÜV (Technischer Überwachungs Verein [English translation: Technical Surveillance Organization]). The use of an independent test laboratory is good safety practice to ensure that an unbiased party comes to the same conclusion as the people who designed and plan to sell the product. This is particularly important in light of the relative complexity of the laser standards. Most individuals, upon first interpreting the standards, find them difficult to navigate. Employing a laser safety consultant often will save time and money in the development of new laser products, as fundamental design decisions can determine the ultimate laser class of the product and thus the market that the product can address.

Eye Safety and Wireless Optical Networks

With the rapid increase in demand for broadband access and the relative expense and delay in deploying fiber, WONs are being viewed as an affordable, reliable, and fast method of providing fiber-like bandwidth. Currently, carrier-grade equipment is being built to provide reliable, redundant access networks. Wireless optical technology proposes to break the last-mile barrier—the relatively short distance between a fiber-optic network and a user. In the case of AirFiber's OptiMesh WON, for example, this is accomplished by Class 1 eye-safe optical transceivers strategically mounted in a mesh configuration on the rooftops of commercial buildings in dense urban areas.

As laser beams become prevalent in public spaces, new challenges arise for manufacturers and users. In response, the IEC is currently drafting a standard under the IEC 60825-1 base standard that specifically addresses open-air communication systems.

Eye-Safe System Architecture

Like many other enabling technologies, potential hazards posed by lasers have been mitigated by the application of safety standards. Lasers are now safely operated in open-air environments as exemplified by the widespread use of laser pointers, laser range-finders, and light shows. Safe operation requires that human access be limited to Class 1 radiation. Evaluating such access for WONs includes the laser-system class, the deployment environment, and beampath. For example, WONs deployed in a space where casual passers-by could potentially intercept a beam requires that the system be Class 1 and incapable of delivering radiation in excess of Class 1 (i.e., under single-fault conditions). This may be an unlikely deployment scenario, because a WON subject to frequent beam interrupts would not be reliable. However, consider beams being aimed into receivers behind office building windows. The space between the terminals may be virtually inaccessible, but mispointing or beam-spread into adjacent windows must be addressed with reliable systems engineering that meets the standard of safety. Similarly, consider beams being sent and received between tower-mounted rooftop terminals. Although the terminals and space between them may be inaccessible, effective beam-stopping and mispointing again must be addressed and solved with engineering solutions.

There is great diversity in the WON system architectures being developed today. Effecting safe public deployment of these systems requires top-level design choices to limit the access to harmful laser radiation. For example, design features such as the number and size of transmitted beams, divergence, power, wavelength, pulse duration, physical size of terminals, and installation location (i.e., network planning) all must be considered in determining laser-system classification and human access. Furthermore, active safety systems that control the laser in response to a potential access event are often implemented. An example is a location monitor (LM) system or an automatic power reduction (APR) system.

An example timing diagram for such an APR is shown in *Figure 2*, and important time scales are labeled. The constraints on the timing are determined by all of the aforementioned factors that are unique to the architecture of WON equipment. The shutdown time of the APR (T_{APR}) refers to the time from the instant an exposure event begins to the time the laser is turned off. This response time is determined such that, under appropriate viewing conditions, a person does not collect greater than the Class 1 AEL in the time T_{APR}. If delays can occur in the system, they must be accounted for and included in the safety analysis. An example may include the time between the onset of the event and the time that the system actually senses it. This example includes the use of a "ping" to test if the link is still blocked. The power of the laser during this pulse and the duration (T_{ping}) are chosen such that the Class 1 AEL is not exceeded. It must be assumed that an observer is still intercepting the beam during this time, and the method to test the link must be inherently safe. Depending on the laser power at shutdown time and the power contained in the ping, a delay (T_{delay}) may be required to remain eye-safe. Lastly, the total time from the event start (t = 0) to the time at which the cleared link is restarted ($T_{restart}$) is similarly chosen so that the system is eye-safe to a viewer who (hypothetically) repeatedly interrupts the beam. A similar Class 1 ping could be used for link-acquisition mode, where two transceivers use a stare-scan method to establish the link during installation, for example.

The AirFiber OptiMesh Class 1 WON

An example of a Class 1–compliant optical wireless system is the AirFiber OptiMesh WON. The product uses rooftop transceiver nodes that communicate with neighboring nodes. The system uses 785 nm lasers transmitting from a two-inch aperture, and employs an APR system that monitors beam access and reduces output power to Class 1 levels in the event of a beam blockage. (The IEC Class 1 AEL for 785 nm is 0.56 mW.) If a beam interrupt occurs, the APR shuts the system down such that a potential viewer does not collect greater than the Class 1 limit and traffic is rerouted to a redundant link in the network. The system then enters recovery mode, where the link is tested to determine if the obstruction is cleared by use of a Class 1 ping. Once the path is clear, the link is re-established for full data traffic. A redundant system of hardware and software was designed such that the system remains Class 1 under single-fault conditions. Thus, it is a CDRH Class I and IEC Class 1 laser system under applicable standards. Independent certification has been performed by TÜV testing laboratories in which test methods as prescribed by the standards were performed by TÜV representatives, including selectively "failing" various subsystems to confirm single-fault tolerance.

In summary, top-level performance requirements included laser eye safety in the design and prototype stages to realize Class 1 operation. The installation environment spans restricted or inaccessible space. However, in the interest of public safety, the system was designed to be Class 1. Rigorous testing was necessary to achieve proper operation, and independent testing was performed to confirm compliance.

Multiple Factors Impact Eye-Safety Classification

Several other wireless optical systems today use 1550 nm wavelength lasers as an alternative to the shorter wavelengths used by AirFiber. These systems are using lasers and components designed for optimal transmission over optical fiber, because this wavelength is the fiber attenuation minimum. A collateral eye-safety benefit is due to the reduced retinal absorption at 1550 nm compared to shorter wave-

FIGURE 2

Example Timing Diagram for an APR Response

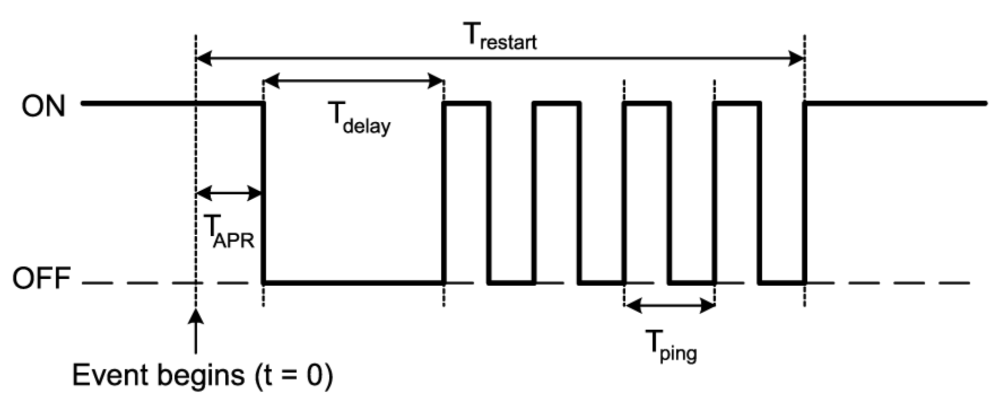

lengths. For example, the IEC Class 1 AEL for 785 nm is 0.56 mW, whereas the Class 1 AEL for 1550 nm is 10 mW—nearly 20 times higher. Does this mean that 1550 nm systems are 20 times more eye-safe? Not necessarily, as the eye-safety and classification of a WON is a function of many system parameters. There are technology trade-offs between long and short wavelengths, such as detector responsivity and active area. In addition, system performance benchmarks such as designed link length, bandwidth, and availability vary. All of the aforementioned determine the amount of laser power required, beam size, and divergence. Thus, to assess the safety hazard of the system, one has to consider the entire system and installation environment.

Safety Standards for WONs

In response to the growth of free-space optical systems for telecommunications, relevant standards bodies such as the IEC and ANSI are developing standards specifically addressing these systems and the unique laser-safety aspects that they present. At the time of this printing, these standards are currently under development by the drafting committees within the standards bodies. It may be possible that they are published in final form in the coming year.

The draft standards at present naturally build from related standards addressing optical fiber systems and outdoor laser light shows, for example. Unique features of WONs, such as LM systems and APRs, are also addressed, as are the environments in which they may be operated and the potential for presenting ocular hazards. For instance, systems beaming lasers near street level, where passers-by may intersect the beam should be rigorously required to be Class 1. On the other hand, systems that send laser beams between towers in restricted environments may be allowed to be in a higher class in recognition of the inaccessibility of the space in which they operate. The committees are composed of laser eye-safety experts, WON manufacturers, optical communication experts, and service providers that will use and deploy such equipment.

Conclusion

With the upcoming deployment of WONs in the public space, laser eye-safety is an important responsibility of the equipment manufacturers and service providers. Laser light poses unique ocular hazards that must be mitigated by inherently eye-safe system architectures, appropriate installation locations, and equipment reliability. The wavelength, power, and aperture of the laser are important, but the entire system and installation must be considered in determination of the relative safety. Current safety standards address laser eye-safety issues and provide guidance to manufacturers and users. However, unique challenges exist in the deployment of these systems in public spaces, and specific standards are being drafted today to address them. The advent of telecommunication-grade WONs and the development of open-air laser-safety standards will enable the delivery of fiber-like bandwidth at a fraction of the cost, while ensuring public safety.

A Class 1–compliant wireless optical network exists today in the AirFiber OptiMesh wireless optical network. Class 1 operation was incorporated into the system by design to meet the most rigorous of eye-safety standards while providing reliable, affordable broadband access technology.

Notes

1. The nanometer (nm) is a unit of measurement used here to describe wavelength (10-9m).
2. The micron (mm) is a unit of measurement used here to describe wavelength (10-6m).
3. The millimeter (mm) is a unit of measurement used here to describe wavelength (10-3m).

Managing Wireless Data Networks in an Era of Change

Gary Barton
Business-Development Manager, OSS Solutions
Agilent Technologies

Eva Thomas
Technical Market Manager, OSS Solutions
Agilent Technologies

Executive Summary

Many service providers see wireless data technology as an opportunity to increase revenue and market-share by providing more valuable services to their customers. But the challenge of adding the core Internet protocol (IP)–based technology necessary to achieve this is much more complex and demanding than many realize. The value proposition has changed from just keeping a physical network up and running, and charging customers for using it, to providing value-added services and content. To succeed, providers will need a unified service and network-management system (NMS)—one flexible and scalable enough to reduce time to market, adapt to managing a more complex environment cost-effectively, and generate a faster and higher return on investment (ROI). Above all, providers must find ways to retain customers with consistently high quality of service (QoS). Optimally, providers should seek flexible operations support system (OSS) solutions that can meet immediate management needs, yet handle the future technologies of the third-generation (3G).

The Outlook for Next-Generation Mobile Networks

Wireless data technology is one of the hottest technologies in the communications industry today. High-speed cable-modem and digital subscriber line (DSL) access have created a hearty appetite for Web-based information, driving business and residential use of the Internet to unprecedented levels. The next logical step is to turn this ubiquitous platform into a vehicle for distributing new kinds of information and services over the airwaves to mobile terminals, such as wireless laptops, cell phones, and other handheld devices.

Cahners In-Stat Group, a market research firm, estimates that 1.3 billion consumers worldwide will have access to wireless Internet service by 2004, with demand for messaging driving this expansion from today's 170 million subscribers. Cahners also predicts that the number of wireless messages sent per month will reach 244 billion by the end of 2004, up more than 800 percent from 3 billion per month in 1999. The study anticipates about 1.5 billion wireless-enabled phones, handheld computers, and other devices in use by 2004.[1]

With consumers' growing desire for access to information anywhere and anytime, we can expect the combination of the Internet and mobile phones to drive up wireless usage significantly throughout the coming years. It is important to note that the business model for such wireless technology is changing to require content as well as access. At the present time, it is unclear what services consumers will value most, and whether these will come from service providers or content providers, which may see mobile devices as just another distribution channel.

Evolving Wireless Data Technologies

Wireless data technologies are evolving on two levels: transmission type and speed, and content. The current digital air-interface technologies—time division multiple access (TDMA), code division multiple access (CDMA), and Global Systems for Mobile Communications (GSM)—have a planned evolution from second-generation (2G) service through packet-based wireless data and on to 3G, core IP–based networks. *Figure 1* shows how these technologies are migrating. How TDMA will evolve is still uncertain, as is the true value of enhanced data for global evolution (EDGE).

While migration plans for the mobile transmission network may be fairly well established, those for formatting and viewing data are still up in the air. Wireless application protocol (WAP), based on an open standard, is currently the most popular viewing format. But the successful deployment of i-mode, a proprietary standard deployed by Japan's NTT DoCoMo, has also caught service providers' attention by attracting 12 million subscribers with its access to numerous Internet sites and electronic mail. These are critical issues for service providers because access to wireless Internet requires a successful viewing scheme.

As service providers and vendors jockey for position and hedge their bets on standards and technologies, consumers are interested in getting quick, cheap access to content, regardless of the technology that delivers it. For years, Internet service providers (ISPs) have provided inexpensive Internet access for wireline companies. Recently, wireless service providers have begun to see revenue opportunity on the content side as a far more promising revenue

FIGURE 1

Wireless Technologies Are Rapidly Evolving to Handle High-Volume Data Traffic

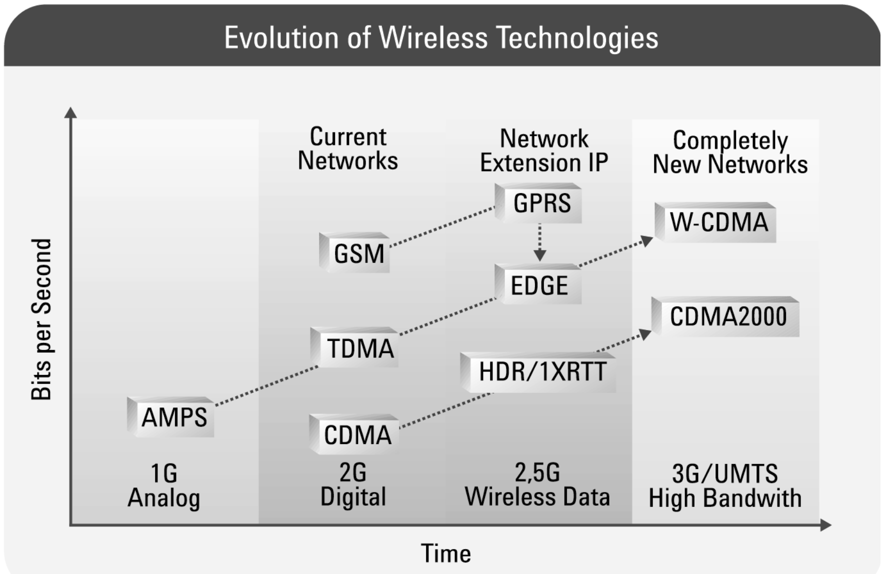

Source: Agilent Technologies

generator than per-minute connection charges—even to the point of becoming wireless application service providers (WASPs) themselves.

Whichever standards and technologies eventually win, innovative wireless data services are clearly on the horizon. To meet growing demand for high-speed data connections, many providers have acquired wireless providers or paid unprecedented amounts for spectrum licenses. Meanwhile, the high cost of spectra and new infrastructure is creating additional pressure for faster time to market, more cost-effective management, and higher QoS and ROI—all of which requires rapid service deployment and tighter operational control.

Wireless Data Players Today

Not surprisingly, the top contenders are current wireless service providers, including mobile cellular providers whose services are based on voice-only technologies. Today's wireless providers can diversify by expanding the type and size of their customer base and increasing customer usage of new kinds of content, either directly, as ISPs or application service providers (ASPs), or through another provider, such as America Online (AOL).

Mobile service providers not only see wireless data as an opportunity to enlarge their reach and revenues, but some consider it necessary for survival. To gain entry into this market, cellular and other wireless providers in Europe have grabbed up highly prized 3G/Universal Mobile Telecommunications System (UMTS) licenses, in some cases grouping together to dilute the cost. Partnerships have been announced between E-Plus Mobilfunk and Mobilkom and between Telefonica and Deutsche Telekom.

NTT's i-mode network is the first wireless data network available to the general public. As a result, the rest of the world is watching to see how its business and technological models work out. Europe, with many general packet radio service (GPRS) trials currently underway, is ahead of other regions, with large-scale rollouts hampered only by the lack of compatible handsets. In the United States, where Internet access is far more affordable and accessible, wireless data lags, with many wireless providers' Web offerings actually being portals to limited content instead of to the full IP–based public Internet.

Challenges and Solutions for Managing Wireless Data

Providers' eagerness to jump into markets while technologies and standards are still evolving is just one indication that they will need great flexibility to evolve their management systems from one technology to the next. For each time a company adds new technology to its mix of services, it must also update its network-management infrastructure, ideally integrating the new with the old. The current challenge is to find a way to transition management systems from voice-centric to packet-based technology.

Wireless service providers' business model is changing profoundly and rapidly. Formerly, their entire revenue stream came from providing a conduit for voice service billed according to minutes used, along with some services such as voice-mail and short message services (SMSs) that are usually billed on a flat-rate basis. With the advent of wireless data networks that offer content, as well as access, providers may choose to give access away and just charge users for content provided.

Those entering the uncertain market for wireless data face some daunting business challenges. To adapt as markets and technologies change—and profit from as yet unspecified opportunities—providers must achieve several key objectives:

- Reduce the time required to get new services to market
- Adapt to managing a more complex environment cost-effectively
- Generate a speedier and higher ROI
- Retain customers with consistently high QoS

It is a tall order, and not everyone is destined to succeed. But the likely winners are those providers that take a more comprehensive, coordinated approach to managing network infrastructure, using their more efficient NMS for competitive advantage.

Objective 1: Bring Services to Market Faster
Providers always stand the best chance of winning customers by being the first to market with new services. Yet as they add network data capabilities and new services based on IP, they must support them with many more network elements (NEs) and element-management systems (EMSs), slowing time to market.

To help mitigate these effects, providers must have the ability to accurately predict future network capacities and perform true end-to-end data and air interface testing.

Implementing a packet data overlay on a cellular network also requires adding widely different vendors and protocols. For example, regardless of which vendor supplies the cellular voice equipment, a GPRS network may involve additional vendors, as in Motorola's solution, which uses both Motorola and Cisco Systems elements.

Adding packet data to cellular networks imposes the additional burden of managing new standards and protocols and OSS software, risking further delays. These, too, require additional expertise, adaptation to installed infrastructure, and extra time and cost for personnel training. The need to integrate and accommodate different elements, vendors, standards, and protocols can cost providers precious time in deploying new services. Wireless data providers must find ways to manage a growing number of diverse network resources more efficiently by simplifying and streamlining their NMSs and reducing the time and cost of training operators to use them.

Solution: A Streamlined, More Efficient Management System
A single, comprehensive management system based on a flexible, OSS framework can provide a number of economies. It allows providers to extend their services and installed networks, integrate various applications to provide a common view of data and operator functions, and automate tasks and processes.

Such an approach enables providers to implement short innovation cycles and add new elements or services without having to upgrade their NMSs or purchase new ones. For example, automating control of the service-delivery process enables the speedy identification and reallocation of available network equipment.

In 3G networks, routers and asynchronous transfer mode (ATM) switches will play a key role in distributing wireless data. The autodiscovery component of an integrated system can search for new simple network management protocol (SNMP) elements and add them automatically to the NMS. Autodiscovery can identify an element's vendor and type, as well as assign default management policies based on that type. Such automation is critical as networks expand rapidly to include many smaller elements, such as routers, which may be added on a daily basis.

A comprehensive management system also provides a single point of control for managing service activation across multiple technologies—a valuable advantage for integrated communications providers (ICPs) and wireless providers with multiple networks. For example, when migrating from one wireless technology to another, a provider may have to support 2.5G and 3G networks simultaneously.

For 3G networks, a broad and robust management system offers the ability to simulate network traffic, measure network capacity and service quality, and have an architecture that is scalable for future networks. The ability to quickly and accurately identify and resolve sources of radio frequency (RF) problems is crucial to compete in today's rapidly evolving wireless market.

Such unified control enables providers to replace multiple system monitors with a common operator interface, reducing the need to train operators on new systems and improving productivity and efficiency by providing visual and operational consistency.

Once such a foundation is established, the provider can use a combination of process automation and workflow to automate the interaction of integrated OSSs efficiently, saving valuable time and reducing errors and costs by capturing, controlling, and automating complex processes (see *Figure 2*).

Objective 2: Cost-Effective Management in a Complex, Rapidly Evolving Wireless Environment
Providing packet data and browser-based content requires wireless companies to manage IP–based equipment—a whole new skill set. They must somehow shift their services and customer base from today's circuit-switched cellular technology to a core IP–based, packet-switched environment. To do this, they need a management system capable of supporting constantly changing networks and of determining proactively when certain NEs or services can be shifted to support new types of customers and their wireless data needs.

Providers need to control costs in spite of manual processes, a substantial increase in the number of resources and separate systems that they must manage, and the need to deliver many different kinds of service. For example, the number of incoming events to the NMS increases exponentially in the IP environment. This produces a similar increase in the number of alerts displayed, threatening to overwhelm network operations center (NOC) staff.

Providers also find that they must incorporate newer standards into their management systems to improve efficiency. Initially, based on American Standard Code for Information Exchange (ASCII) or common management information protocol (CMIP) standards, wireless network interfaces are evolving toward the use of SNMP and common object request broker architecture (CORBA), both of which can help reduce costs. For example, SNMP–based equipment, which predominates in packet networks, can

FIGURE 2

Workflow and Process Automation Can Be Used to Manage Complex Processes from End to End, with Minimal Human Intervention

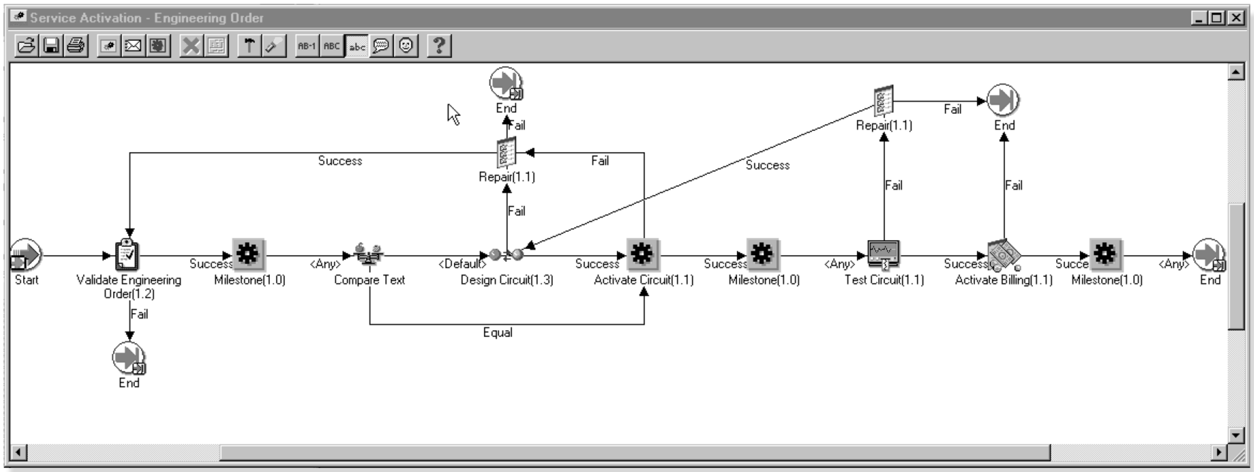

Source: Agilent Technologies

be adjusted to monitor for signs of network overload and thus prevent outages. The CORBA standard is now being used to assure compatibility between EMSs and NMSs and to link OSSs.

Solution: A Highly Flexible, Open and Interactive Management System

To keep pace with rapid change and growth, wireless data providers will need an open management platform with integrated applications that enable not just passive monitoring, but also interaction with the network—even to the point of anticipating and resolving problems before customers are aware of them.

Policy-based management of network data, thresholding of performance data, and process automation all have a useful role here. When combined as packaged fault management policies—sets of reusable rules that can be applied to different groups of elements—such techniques can help providers to get their management systems up and running rapidly. The additional option to design and implement their own management policies lets providers customize the way their management system responds to various problems.

Automated techniques—such as thresholding, correlation, suppression, and escalation—can dramatically reduce the number of alerts shown and point to the most important events, or root causes, saving valuable staff time and effort.

For example, performance thresholding inhibits the generation of an event until a specified number of alerts has occurred within a certain time interval. NOC staff need only respond when the assigned severity exceeds a predetermined level within that interval. Fault correlation enables the system to look for a pattern of alerts and generate a new alert based on their actual cause. Alert suppression limits the display of alerts to just the highest level of alert and suppresses sympathetic alerts. Incoming alerts can also generate trouble tickets automatically and trigger other automated test routines and related processes.

Network operations also become more cost-effective when a provider can normalize network alerts and performance data from all network resources and view them from a common alert display. This requires the tight integration of essential applications for managing network configuration, faults, performance, and accounting, making such data centrally available and visually easy to interpret.

In addition, because the infrastructure of a 3G network is spread over a wide area, operators can benefit from geographical displays that provide an instant, accurate view of network problems and adversely affected areas (see *Figure 3*). This provides operators with a picture of how traffic is flowing to see where loads are high, an important benefit because the virtual—and as a result, largely self—healing nature of a 3G network makes it difficult to locate network problems.

Once key management applications are integrated, a process-management system that combines workflow and process automation can be used to control the end-to-end management of tasks, whether automated or manual. This increases the management system's scalability to handle ever-larger volumes of traffic and data.

For example, the entire service activation process can be managed from end to end, including validating work orders, provisioning NEs, testing services, setting up service-level agreement (SLA) monitoring, and activating the billing system.

Objective 3: Achieve a Better ROI

Given the high costs of wireless spectrum licenses, additional NEs, and managing increasingly complex processes, service providers are looking for ways to get a faster return on their investment. Because providers must use their spectrum licenses within government-established time frames or lose them, many providers are under pressure to meet very tight deadlines for getting 3G networks up and running. Others need to do so just to recoup their expenditures. The financial market is unforgiving, and any provider that does not implement services rapidly, with good prospects

FIGURE 3

Map-Like Displays Indicate the Location of Network Problems

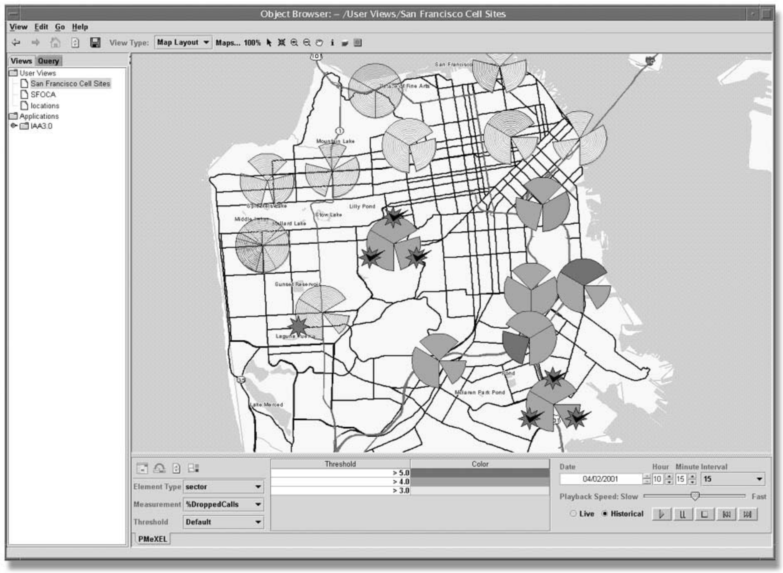

Source: Agilent Technologies

for satisfied customers and high returns, will find it difficult to finance additional expansion.

To attain higher ROI, providers must increase network availability and find ways to charge for the value of content, as well as for connection time. The problem with network buildout is that no one can actually predict demand for wireless data. For example, no one knows how much bandwidth will be used for data, how the use of voice and data may shift during the day, or what the overall traffic pattern will look like. Given such uncertainties, providers may suddenly need to expand their network or face bandwidth shortages.

Another revenue stream for service providers is in the ability to bill other service providers for the use of the signaling system 7 (SS7 network), trunk network, or service database. Equally as important is the ability to verify the charges billed by other service providers for the use of their networks and services.

One chief obstacle to expanding the mobile voice revenue model to include packet-based content is providers' inability to capture—or even recognize—the limitless possibilities, let alone bill for them. For example, a smart handset equipped with a global positioning system (GPS) chipset can already alert a subscriber to an offer of interest at a nearby store. Providers will eventually want a way to track such instant messaging and bill the merchant involved.

Switching from flat-fee and time-based billing to charging customers for the value of content requires more sophisticated systems than currently available. Providers need more direct ways of tracking data usage than current systems allow, to enable advanced billing based on content value, as well as call duration. For while subscribers may be willing to pay premiums for the transmission of critical wireless data during peak usage hours, providers need the ability to distinguish such use from less expensive service, as well as to determine peak usage by traffic volume and time of day.

Solution: Advanced Management Software for Performance and Billing

Advanced management software that ensures higher network availability and optimizes the value of existing equipment can help providers to achieve a higher rate of return on their investment in network infrastructure and spectrum licenses. Using such an approach, wireless data providers can implement new pricing policies based on content, as well as time, and reduce or at least postpone additional capital asset expenditures, allowing them to keep more revenues and generate higher levels of revenue per subscriber.

Congestion in the radio access and IP transport network, due to capacity constraints and equipment problems, reduce not only subscriber access to network time and content, but also revenue for service and content providers.

In fast-changing wireless environments, providers will need to monitor system capacity and utilization continuously, analyzing usage trends and forecasting network buildout requirements. An easy-to-use performance-management system can dynamically manage large volumes of performance and traffic data and be used to forecast performance trends to avoid disappointing customers with poor QoS (see *Figure 4*).

In this way, advanced performance management can help providers to keep more of their revenues by assuring network availability, avoiding service degradations and maintaining high call- or connection-completion rates, preventing slowdowns in data throughput and excessive dropped calls and busy signals.

FIGURE 4

Graph-Like Displays Track Network Performance

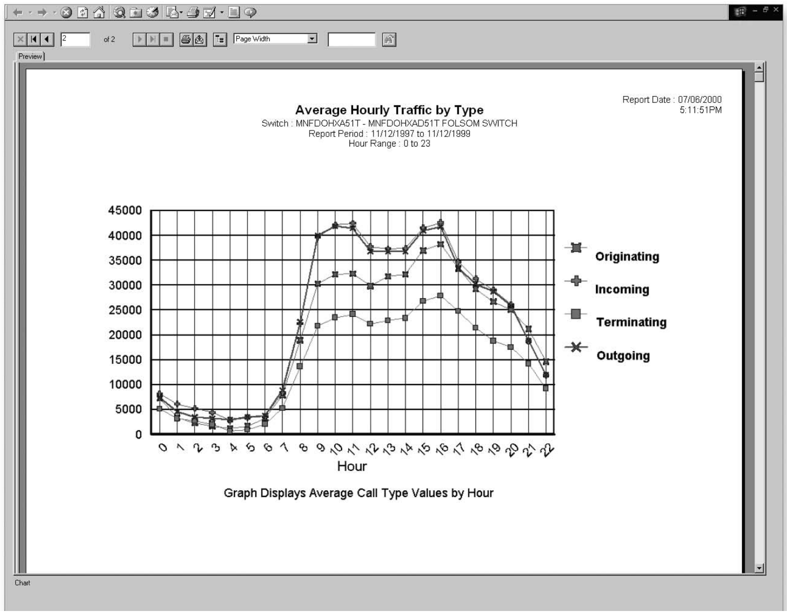

Source: Agilent Technologies

The accurate collection and analysis of performance data—along with the ability to share the results in reports—is critical to responding immediately to network problems that may affect customers and to understanding where resources can be redeployed or adjusted for optimum benefit and the highest ROI.

Real-time performance management enables providers to minimize expensive performance problems by establishing network resource threshold values that, when surpassed, reveal problems that require immediate attention. It also enables strategic network planning based on the use of performance data feedback and analysis to ensure the correct sizing of network buildout, avoiding excess costs.

The promise of a wide range of new services gives providers the opportunity to implement new pricing policies based on content rather than just time. But to recognize this new revenue stream, providers' accounting mediation and management systems must be able to make billing records from multivendor sources compatible with an upstream billing system.

Billing records for 3G services will contain more call attributes than ever, including details of service priority and quality and of content provided. Useful for more than just billing, these call detail records (CDRs) also help to reduce churn and capture more revenue through their use in QoS analysis, SLA management, and applications that perform secondary functions, such as fraud analysis.

CDRs—a direct record of a customer's experience—currently provide a limited amount of information. In response, some infrastructure vendors are adding charging gateways to their products, making it easier for providers to collect detailed usage records. Many are also working to create the IP data record (IPDR), a new standard for formatting records from IP elements that includes information relevant to billing, such as packet count, time of day, delivery options, service-quality level, and so on.[2]

Objective 4: Retain Customers by Preserving QoS

In an era of tough competition, customer satisfaction has become wireless providers' top priority. Heavy increases in IP traffic threaten to degrade both network capacity and performance, affecting perceived service quality. By taking on wireless data services, providers face even larger challenges with customer churn, because services carried over the data stream will be more sensitive to errors than voice.

For while a dropped voice call is annoying, a broken wireless data connection in the midst of a stock trade is a threat to a customer's revenue stream. Once wireless networks are based on IP, individual packets of wireless data may follow different paths to their destinations, and service providers may not realize there is a problem until a customer calls to complain.

In 3G systems, some services will be far more vulnerable to QoS problems than others. For example, short-term fluctuations such as latency and jitter will not affect Internet access but will definitely degrade the quality of voice (e.g., voice over IP [VoIP]) and streaming applications (e.g., videoconferencing and teleconferencing).[3]

Measuring these fluctuations will be especially critical to assuring QoS in services that involve streaming packets. For while an IP network is by nature self-healing and can retransmit packets that are lost or have errors, such delays can significantly decrease total network throughput.

Monitoring and controlling a wireless data network for QoS will be much more difficult. For instance, while GPRS networks are actually two separate networks—voice and data—they must be managed as one. Each network's performance affects the other because they share capacity constraints and common elements such as cell sites and home location registers (HLRs) (see *Figure 5*).

Failure in the core GSM network could cripple the data network's ability to route packets. To support higher QoS at the network level, providers will need management systems that can pinpoint the location of specific problems and automatically take action to correct them.

To maintain QoS at the business level, service providers will have to satisfy both content providers and their customers, who expect high-quality, uninterrupted service. They need the ability to monitor and manage QoS from the customer's perspective. To maintain customer loyalty and reduce customer churn, service providers must monitor not only the customer experience as it pertains to coverage and penetration, but also measure and monitor the response times and quality of the services that run across the NEs.

Wireless customers already expect the same quality with wireless data that they get with tethered Internet access. To compound the problem, many providers anxious to build market-share for 3G and beyond have set unrealistic expectations for QoS that may later cost them both market-share and revenue. When 3G becomes a reality, providers will have to deliver on their promises of untethered, anytime, anywhere high-speed wireless access to the public Internet—or suffer the consequences.

For reasons like these, SLAs, once unknown in the wireless industry, are increasingly popular with content providers and business users that want guarantees of both access to customers and a specified QoS—with rebates to compensate them for receiving anything less. For if a service provider's data network goes down, the content provider stands to lose both revenues and customers.

Solution: A Robust Management System to Maintain High QoS

QoS is a relative measurement of the service a customer experiences—a combination of many different factors that contribute to customer satisfaction or dissatisfaction. To monitor QoS means monitoring faults, performance, and accounting (CDRs). For an adequate understanding of the service being supplied through the network, providers must even perform simulated test calls. They need a management system that can collect, correlate, and analyze this information for an accurate picture of QoS.

To continue to attract and retain customers and guarantee service levels, wireless data providers need a robust management system that can provide tight end-to-end monitoring of connections and services to assure that customer service levels are met. Unified fault and performance management can play a large role in assuring the QoS across the entire network. A fault-management system can quickly alert operators to problems without overwhelming them with unnecessary alerts, which slow down analysis and response times. Real-time thresholding in performance management can warn of impending problems, including capacity overload.

Service providers must have the ability to verify the quality of network coverage. A management system that provides both in-building and outdoor signal strength measurement allows for proactive measures when addressing current customer need and assessing future customer demand.

Such a system can also provide the automation of systems and processes needed to support high QoS in such a complex, changing environment. The more reliably a provider can keep the network up and running, the less the company's reputation for service will suffer and the less it will have to pay customers in SLA credits. An automated reprovisioning response to an equipment outage can support an SLA far better than a slower, error-prone manual dispatch.

FIGURE 5

GPRS Voice and Data Networks Are Vulnerable to Failure Because They Share Common Elements and Capacity Constraints

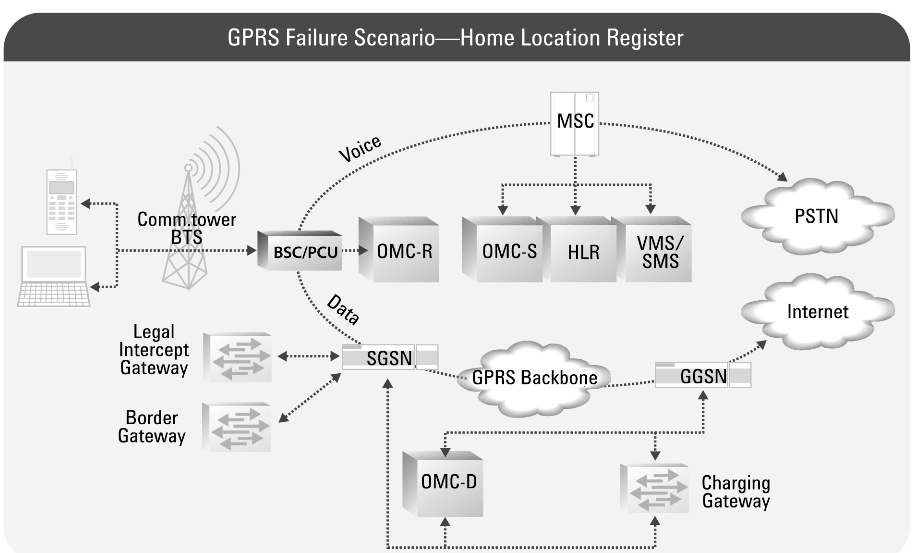

Source: Agilent Technologies

Sophisticated forecasting and analysis, using data available from many different management processes, can also play an important part in enabling a provider to identify new opportunities and plan for future services—a necessity in continuing to attract and retain customers. Forecasting can indicate wireless network usage trends and reveal where and when additional capacity will be needed.

Appendix: The Evolution of Wireless Technology

As cellular technology advances, the concept of transmitting high-speed, multimedia wireless data is approaching reality. Over time, wireless technology has evolved from first-generation analog voice networks with limited capacity to more powerful second-generation (2G) digital cellular voice networks based on TDMA, CDMA, and GSM technologies. As demand for wireless data has risen, 2.5G technology was developed to supplement 2G networks with separate GPRS or 1XRTT (wireless data for CDMA) networks—typically packet-based, but not full IP.

TDMA provides a higher density of calls, better sound quality, and improved security. CDMA makes better use of the frequency spectrum and handles even more calls. Some organizations have used high-speed circuit-switched and cellular digital packet data (CDPD) technologies to allow mobile units of police, fire, and utility personnel to send and receive critical information. Recently, several carriers began using CDPD to provide access to business information and applications on a Web-like proprietary wireless network as an intermediate step to full wireless access to the Internet.

For example, Sprint's Wireless Web for Business service boosted transmission speed from 14.4 kilobits per second (kbps) to 56 kbps and gave users access to e-mail, as well as Microsoft Exchange and Lotus Development's Notes. However, such services cannot match the speed and rich graphics of full-screen color Internet displays, so customer demand will continue to build for the real thing.

2.5G Technologies

GPRS, 1XRTT, and EDGE represent higher-speed transitions to 3G. GPRS is a packet-based technology designed to add data to GSM technology for high-speed wireless data service. Currently used for wireless Internet browsing, e-mail, and various push technologies, GPRS can reach speeds of 115 kbps. Mostly in the trial stage today, GPRS has limited implementation by companies such as BT Cellnet.

EDGE technology marks the convergence of TDMA technology with GSM networks to enable wireless data services. GSM networks will initially deploy networks based on GPRS, then may migrate to EDGE if it becomes widely available. TDMA will skip the GPRS phase and go straight to EDGE or Enhanced GPRS (EGPRS-136). Because the transition from GPRS to EDGE only involves changing the air interface technology, it will mainly require making software changes.

GPRS and EDGE will require a new subsystem for carrying packetized data. A packet control unit at the base-station controller (BSC) will split off the packetized data and send it down to the GPRS network, where it can connect to an Internet portal. The GPRS subnetwork will add elements that are new to cellular systems but standard for IP–based networks. Most of the new equipment will consist of ATM switches and routers, with additional standard UNIX workstations for specialized software and SNMP interfaces.

CDMA Wireless Data

CDMA technology is taking a separate path from TDMA and GSM. Whereas TDMA and GSM add a separate network for handling data calls, the CDMA network is more integrated. 1XRTT, the first version of high-bandwidth data for CDMA systems, will convert voice transmissions to data packets within the network to make more efficient use of the spectra than GSM can provide. As the CDMA data standard develops, it will move to 3XRTT for even higher data transmission. Whereas current GPRS networks must establish a set amount of capacity for voice and data, 3XRTT can dynamically allocate capacity depending on need.

3G

Known in Europe as UMTS, 3G networks are designed to accommodate high-speed multimedia data applications in fixed broadband and mobile wireless environments. The original idea was to establish a single air interface on a single frequency for worldwide roaming, using a standard handset, but the abundance of frequency and use of multiple technologies in the United States won out.

The selection of CDMA for its higher bandwidth than other technologies did not prevent the proposal of competing variations of wideband CDMA (WCDMA) technology. Consequently, three different air interfaces—WCDMA, CDMA2000, and UMTS terrestrial radio access–time duplex division (UTRA–TDD)—were approved for the final edition of the IMT–2000 standard. Countries vary in the frequency spectrum available to them, leading to the approval of several frequency ranges, as well.

The core 3G network will evolve from the equivalent of 2.5G core network systems. A new component of network infrastructure for 3G networks is the media gateway, which resides at the boundary of different networks to process end-user data, such as voice coding and decoding, to convert protocols and to map the QoS.

In addition to requiring a new base-station system, 3G networks will require more base stations than comparable 2G networks. 3G networks that combine cellular voice and packet data technologies will also require new radio and core NEs.

Notes

1. Cahners In-Stat Group, "Wireless Data/Internet Market Gains Momentum: Five-Year Subscriber Forecast," No. MD0003MD, August 2000, p. 41.
2. The goal of the IPDR initiative is to define the essential elements and format for data exchange among network devices, OSSs, and business support systems (BSSs) to enable next-generation IP networks to operate efficiently and cost-effectively.
3. Latency refers to the delay of packets traveling through the network, while jitter refers to packets that are out of sequence.

Quality of Service for Wireless Networks

Selim Baygin
Product Specialist
SOTAS, Inc.

Mario Margolis
Product Manager
SOTAS, Inc.

Introduction

The purpose of this paper is to address the increasing need to monitor quality of service (QoS) over wireless networks driven by the high penetration of voice portals.

This paper will cover four main topics:

- Current wireless market situation
- A case study of one company's application note and test results
- Benchmarking versus continuous monitoring methods
- A cost-effective monitoring strategy

Current Wireless Market Situation

A current lack in reliable wireless communications is impeding wireless technologies from becoming widely accepted, especially for business-critical applications. Declining prices in wireless services do not present a sufficient incentive for customers to substitute their current wireline communication with wireless communication. Delays in adoption of 2.5G and 3G technologies are also stalling the penetration of data applications. Furthermore, voice-driven menu portals are stepping up as an alternative solution in today's competitive marketplace, requiring higher levels of voice quality and network availability.

Benchmarking QoS twice a year is an extremely costly procedure for wireless providers today and does not provide a true representation of the QoS the customer receives on a day-by-day basis. In fact, it seems that QoS dynamically changes over time based on location and time of day. Continuous monitoring from an end-to-end perspective, then, is needed. Such continuous monitoring will allow carriers to address and troubleshoot quality issues on a proactive basis, thereby enhancing their abilities to reduce churn, to increase the amount of billable time their subscribers incur, and to differentiate their services in the marketplace based on quality and not price alone.

Background

Wireless
Wireless communications, with the right focus on quality, could become a predominant means of communication in a developed society, provided that quality improves to near–toll level for both availability and reliability. Indeed, estimates, such as those provided by ResearchPortal.com (see *Figure 1*), project substantial increases in the number of both data-centric as well as voice-centric small form-factor (SFF) devices.

Nearly half the U.S. population owns a wireless phone, but only 3 percent of those users are relying solely on that medium for their day-to-day communication needs. This is indicative of the perceived quality and reliability of wireless networks by the customer. Though prices keep dropping, approaching toll-level structures, wireless QoS has not been able to improve at the same rate.

According to Consumer Reports, more than half of customer complaints are related to dropped calls. JD Power and Associates' wireless customer satisfaction index 2000 indicates that call quality (30 percent) has the highest impact on consumer satisfaction, followed by pricing (22 percent), customer service (20 percent), cost of roaming (3 percent), handset (2 percent), billing (2 percent), and others (21 percent).

Voice Portals
Voice portals are emerging as a suitable alternative to slow wireless data-transmission rates. The Kelsey Group estimates that 18 million consumers will use a speech recognition portal by 2005. By then, revenues for these portals will top $12 billion.

These portals serve as a stepping stone to high-speed data services that 2.5G and third generation (3G) technologies are promising by 2003. High levels of voice quality, however, are essential for these portals to become widely utilized in a wireless environment.

Case Study: Introduction

As carriers start to focus more keenly on delivering seamless wireless service to the marketplace, the sheer number of vendors and the technologies they use to construct their networks begs the question, "Is the customer is getting the best service in terms of quality and availability?" It is the premise behind this simple question that can spell the difference between capturing a large portion of the market—becoming the leader—or capturing the remaining portion of the market—becoming a follower.

In light of the aforementioned facts, one company, call it Company X, has attempted to apply its knowledge in the end-to-end QoS arena to the wireless marketplace, proposing to implement a cost-effective continuous-monitoring strategy.

Company X began its project by looking at the various factors that would have an effect on the quality and availability of wireless service in major metropolitan areas. The initial stage of this project focused more on the applicability of the conventional QoS metrics within the wireless world. Since this was uncharted territory, it was interesting to note the effects certain factors, such as location and time of day, had on quality and availability metrics.

The next section provides a brief overview of Company X's system setup. It is important to keep in mind here that the sample size of this test campaign was limited and that the main idea was to look at the various relations between QoS metrics and environmental and temporal factors.

Following the system setup information, some key findings, which may provide valuable insights to end-users and carriers alike, are presented. It must be noted, however, that these results are not conclusive and should not be viewed as a general rule in the wireless communications world.

Case Study: System Setup

To conduct a successful wireless test, it is critical to interface to the wireless network in the most seamless method possible, simulating the day-to-day usage scenario of a regular customer. While this may seem like a trivial task, dedicated research in this area led to the selection of several wireless phone interface vendors that would deliver the required functionality.

Two CallWave™ units were used in conducting the tests—one of them connected to the wireless network using a proprietary interface from Vox2 for a digital Nokia 5160 phone, and the other connected to the public switched telephone network (PSTN). Both connections were made over the two-wire registered jack–11 (RJ–11) ports on the CallWave units.

Case Study: Test Setup 1

The main objective in the design of the wireless application was to keep one unit fixed at Company X's headquarters in Gaithersburg, Maryland, and perform tests from the mobile CallWave unit to the stationary unit. Using this setup, a controlled experiment environment was created, which in turn ensured a more reliable and scientific test scenario.

On the mobile-unit side, three locations were chosen for this test. The first location was company headquarters in Gaithersburg, which simulated a professional on the road from work to home. The second location was in the well-populated Georgetown area of Washington, D.C., and was supposed to simulate an end-user accessing the wireless network from a typical metropolitan location. The third location was Arlington, Virginia, which simulated a boundary metropolitan-area subscriber.

Time was the other dimension in this test campaign. The testers wanted to determine whether placing calls at different times of the day had any significant effect on the quality of the calls and the availability of the network. Calls were placed during both busy hours and non–busy hours, covering a 24-hour period with 30-minute intervals between each test call.

Initial tests included a limited number of measurements in order to keep the test setup simple and easy to interpret.

FIGURE 1

Worldwide SFF Device Forecast

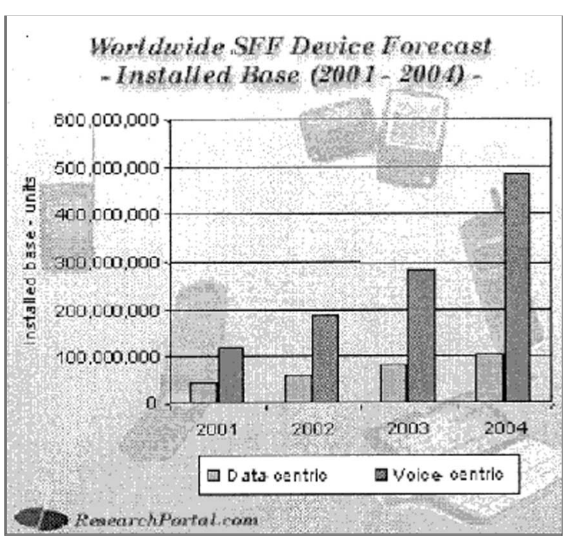

FIGURE 2

System Setup—Interfacing to the Wireless Network

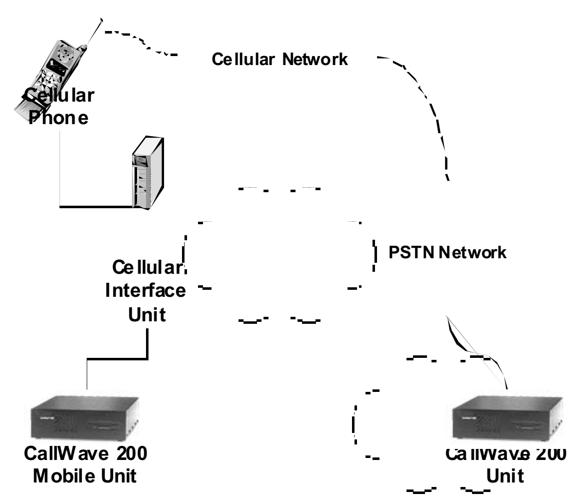

The following key performance indicators were used in the test setup:

- Call-completion rate (CCR)
- Post-dial delay (PDD)
- Perceptual analysis measurement system (PAMS)

Case Study: Test 1 Results

This section discusses key results in the following three areas:

- Network availability
- PDD
- PAMS—listening effort/listening quality

For each area, two perspectives are given: results by location and results by time of day.

Availability

The main measure of network performance is availability. To measure availability, the completed calls and dropped calls were analyzed.

Figures 3–6 illustrate a clear difference between how calls are handled in Gaithersburg and how they are handled in Arlington and D.C. Call completion was significantly higher in Arlington (where CCR was at 80.48 percent) and D.C. (78.65 percent) than it was in Gaithersburg (58.66 percent). The difference in call completion between Arlington and D.C. was not significant.

Both Arlington and D.C. show declines in CCR between the hours of 1:00 pm and 9:00 pm. Gaithersburg shows a similar decline, but it is not sustained due to the variability of the results.

FIGURE 3

Call Disposition—Gaithersberg

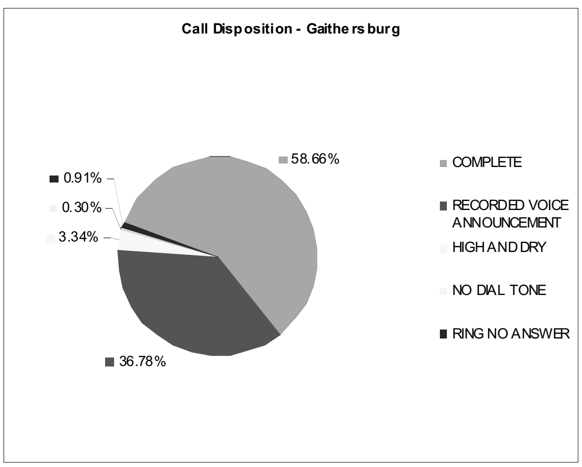

FIGURE 4

Call Disposition—Arlington

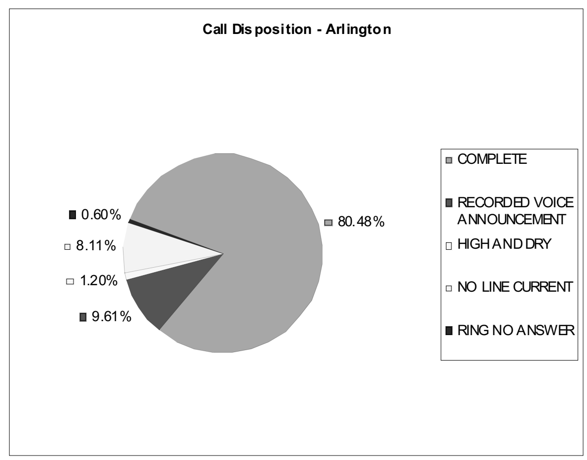

FIGURE 5

Call Disposition—D.C.

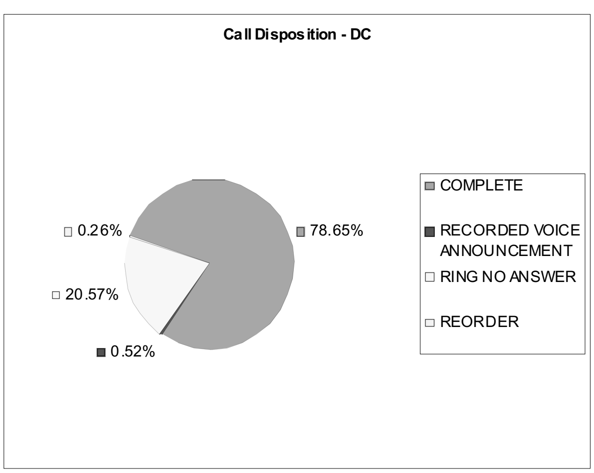

FIGURE 6

Average Call—Completion Ratio by Hour

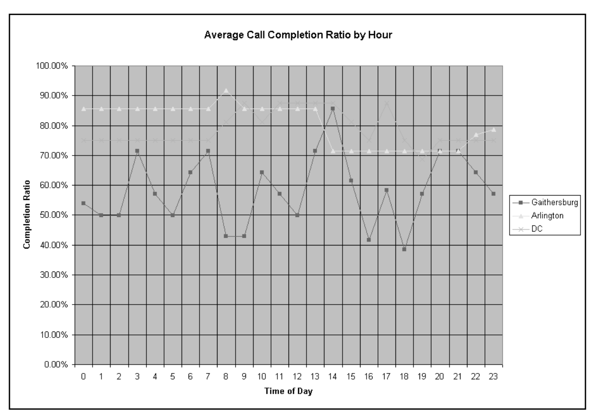

PDD

Another important performance indicator that was taken into consideration in this test setup was PDD (see *Figures 7–8*). PDD indicates how responsive a network is in handling call requests by end users.

Looking at test calls that were placed over a 24-hour period covering seven consecutive days, the conclusions can be drawn from *Figure 8*:

- The highest PDD was observed in Gaithersburg (10.54 seconds).
- Arlington and D.C. had the lower PDD observations (7.74 seconds and 8.31 seconds respectively).
- On an hourly basis, there is severe variability in the results; but Arlington experienced the highest PDD measurements at 2:00 p.m., which coincides with its lowest performance in completion rates.

PAMS

The last measure of QoS here was PAMS. Listening effort and listening quality were evaluated using scores ranging from a low of 1 to a high of 5 on a linear scale. The purpose of using PAMS scores is to understand how perceived quality of voice changes by location and over a period of hours. *Figure 9* illustrates a comparison by location:

Figure 10 shows that calls placed from Arlington and D.C. received slightly higher scores on both the listening-quality and the listening-effort measurements than those calls placed from Gaithersburg.

The higher results for Gaithersburg's observations make the measurement on par with those of both Arlington and D.C.

On an hourly analysis, we find that the listening-quality results from Arlington and D.C. are very consistent, while results perceived from Gaithersburg are lowest between the hours of 3:00 pm and 7:00 pm.

Time-of-day analysis does not reveal any significant results on data collected from any of the locations.

Case Study: Test Setup 2

For this test, a credit-card voice-recognition portal was chosen to perform test calls from a CallWave unit. A recording file that would spell out the credit card number was played

FIGURE 7

Average PDD Comparison

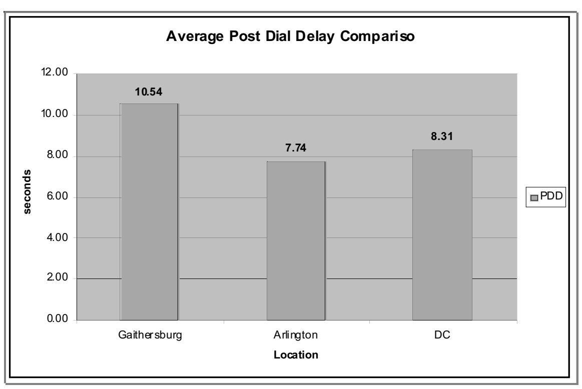

FIGURE 8

Average PDD by Hour

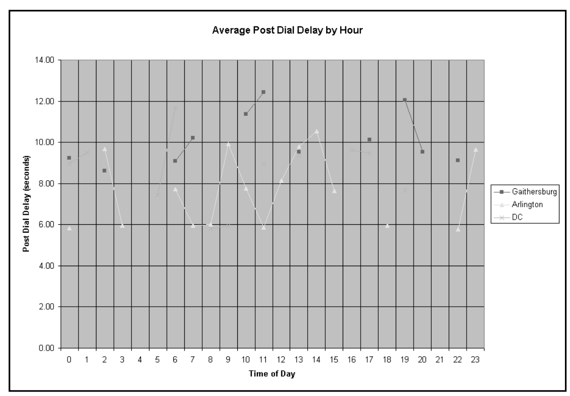

FIGURE 9

PAMS Scores

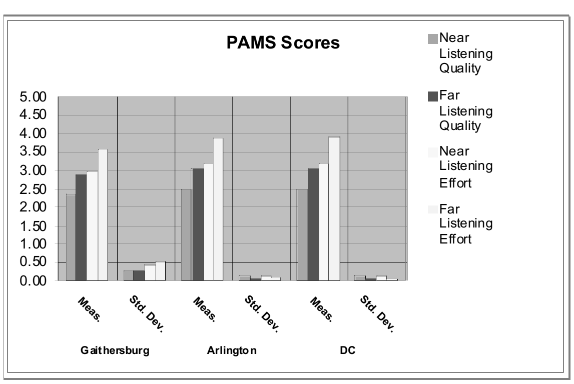

FIGURE 10

Average Listening Quality by Hour

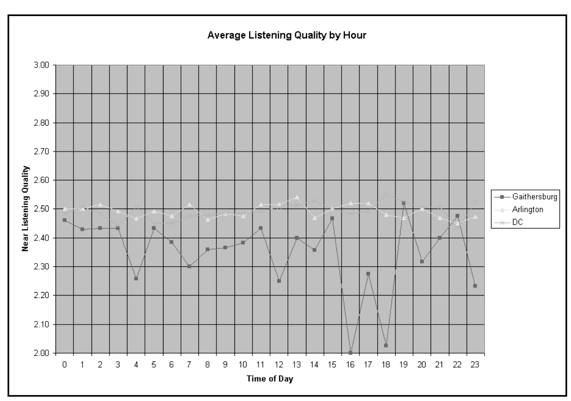

from the CallWave unit. If the portal recognized the number, it would provide the current credit balance. If not, it would ask the user to "say" the number again.

Twenty calls were placed through the cellular interface, and twenty calls through a landline. The calls were placed every two minutes.

The following results, which are graphed in *Figure 11*, were obtained:

Cellular

Attempts: 20
Completed to the platform: 14
Dropped: 6
Completion ratio: 70%
Recognized by platform: 3
Recognition ratio: 21%

Landline

Attempts : 20
Completed to the platform: 19
High and Dry: 1
Completion ratio: 95%
Recognized by platform:13
Recognition ratio: 68%

Next Steps

1. Gather a significant sample size.
2. Correlate recognition ratio with listening-quality (PAMS) results to arrive at acceptable levels.
3. Incorporate similar setup over IP to determine PAMS acceptable levels, and evaluate dual-tone mutlifrequency (DTMF) recognition.

Benchmarking once or twice a year only covers the first two steps required in effective monitoring strategies. Network conditions are constantly changing with the implementation of new services, carriers, and technologies. Prices keep dropping, and QoS expectations keep increasing. Carriers need to compete and be able to differentiate their services using something other than price.

Through continuous monitoring, carriers can always understand the current QoS that they are offering as perceived by the end users and understand how it compares to their competitors' QoS. They can, therefore, be proactive in addressing QoS issues before their customers even know they exist, which can, in turn, help to reduce churn. Isolation and troubleshooting of QoS problems can only be done in a proactive way with continuous QoS monitoring and network element correlation. Benchmarking alone will not provide adequate measurement. Current benchmarking methodologies—conducted only annually or semi-annually—are price prohibitive and make it impossible for carriers to engage in continuous monitoring.

Another major issue in wireless communications today has been the relative shortness of call duration when compared to landline calls. This has been due to both price and QoS. Thus, improving QoS will allow carriers to increase CCRs and average conversation time—two desirable goals from a carrier's perspective since more completed calls mean more calls a carrier can bill, and since longer conversations mean higher bills.

Proposed Cost-Effective Monitoring Strategy

This section describes a cost-effective methodology that carriers can implement in a short period of time.

Establish Baseline Network Characteristics

These will allow wireless service providers to identify and understand the current QoS that they are providing as perceived by their customers. This baseline can be used later as a reference when comparing and differentiating their services with their competitors' services.

Set Benchmarking and Alerting Thresholds Based on End-To-End Subjective Measurements

"Subjective measurements" here refers to mean opinion scores (MOS), call clarity indexes (CCI), PAMS, etc. that provide a perceptual measure of what is acceptable to the customer and what is not.

These thresholds can be static—standards, design, competitive information, service-level agreements (SLA), customer expectations, regulations, and financial targets—or

FIGURE 11

Percentage of Calls Recognized by the Voice Portal

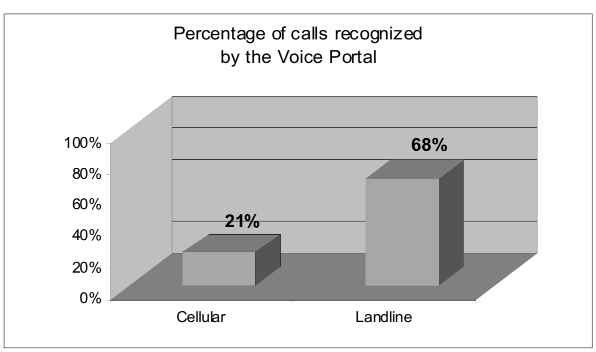

FIGURE 12

Benchmarking versus Continuous-Monitoring Methods

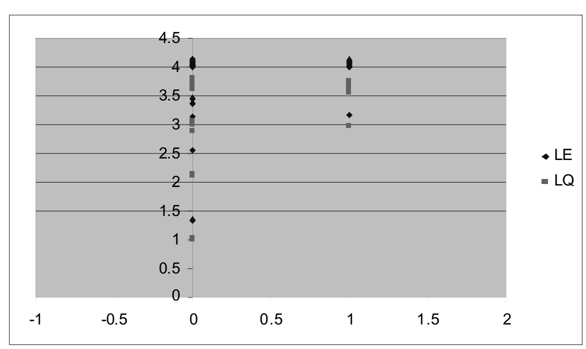

dynamic—profiling systems that determine what is normal behavior and what is not based on "learned" QoS from specific locations, carriers, customers, and destinations).

Prioritize Alerts Based on Financial Penalties and Response Times

These alerts are usually specified in SLAs, although they can be specified based on the importance of services, region, customer, etc.

Subjective measurements need to be correlated with objective measurements and network elements (NE) for problem isolation and troubleshooting

Once thresholds are exceeded, a monitoring system needs powerful correlation in order to determine what the cause (root-cause analysis). In this case, subjective impairment information needs to be correlated with detailed objective impairment information, global positioning system (GPS) information, radio frequency (RF) information, call detail record (CDR) information, and physical-layer–availability information in order to get a full picture of what caused the threshold to be exceeded.

Continuously Monitor To Address Ever-Changing Conditions

The process needs to be constantly active. This is the only way that carriers are really going to be proactive and detect and address any deterioration in QoS before their customers become aware that such deterioration exists. Important business decisions are made constantly, not once or twice a year. Such constancy should similarly apply in maintaining levels of quality with network conditions that are constantly changing.

Provide Reliable and Timely Information Customized for Each Functional Area of the Organization

An effective monitoring system is capable of not only measuring the above characteristics but also disseminating the data, converting it into the right information at the right time to the right decision makers.

Conclusion

Several conclusions can be drawn from this test scenario:

1. Location and time of day *did* have an effect on network availability (CCR).
2. Location *did* have an effect on PDD and PAMS.
3. After a period of time, one will know if results change over time.
4. IIsolating and troubleshooting problems now require more significant sample size; mobile solution to evaluate dropped calls due to handoffs; and correlation with RF levels, geographical location, and switch CDRs.
5. Continuous monitoring, especially in such a high-churn environment, should be implemented to allow provider to become proactive instead of reactive and to increase the amount of billable minutes.
6. Calls over the wireless network *did* have a considerable impact on voice-driven applications that rely on speech-recognition technologies.

7. It is difficult to set acceptable levels, especially in an environment where customer's QoS expectations are believed to be lower.
8. Impairments with the highest correlation-to-recognition ratio must be found; if not PAMS, others (SG12—58 PSQM Noise, PSQM+ clipping and distortion PESQ) should be found.

The next steps are to conduct more geographically dispersed tests. These will involve more locations as well as moving targets to analyze the effects that cell handoffs have on availability, reliability, and suitability of the wireless network. This approach will also provide a closer approximation to a true customer experience. The above test results already indicate poor quality performance of a wireless network, even though the conditions of the controlled environment test were almost ideal. A more dramatic impact can be expected when the suggested true-customer-experience testing scenarios are entered.

Mobile Operators: The Sky Isn't Falling

Getting Back to Business Basics Will Get the Industry Past the Current Cash Crunch

Brian D. Bolliger

Director of Global Mobility Marketing and Strategy, Mobility Solutions Group

Lucent Technologies

There's been much hand-wringing of late over the costs incurred by service providers around the globe in securing spectrum for new-fangled high-speed mobile Internet services. So much so, that in the latter half of 2001, many deployments had been postponed or scaled back. Among reasons cited were technical hurdles, doubts about the business model for next-generation services, skittishness about further capital investment, and the availability of next-generation devices.

We have been here before. Anyone with the benefit of longevity will recall similar conditions when the concept of commercializing cellular communications was first broached in the 1970s. Regulators were initially fearful about having fixed-line ratepayers subsidize what appeared to be a luxury for rich business people. Then, when analog mobile services became wildly successful, similar doubts were raised in the 1980s about the cost and viability of second-generation digital personal communications services (PCS) to be offered in new spectrum by both incumbents and new operators.

To anyone who has been around the industry for the past three decades, the current concerns being raised today about third-generation (3G) networks ring familiar:

- The rates that operators have to charge to recoup their investment will be prohibitive, and so the services promised by mobile providers will appeal only to a select group of rich people or "road warriors."
- You can't surf the Net from a mobile phone like you can on a desktop computer (updated version of the argument that cellular can't match the voice quality and reliability of landline telephones, so why bother?).
- New upstart operators will be unable to compete against embedded incumbents who have a head start and already have broad-based coverage.
- Wideband code division multiple access (CDMA) technology is radically different from second generation (2G), hence it is best to continue to invest in 2G and 2.5G until 3G is perfected (an updated version of the 1980s argument that spread spectrum CDMA is unstable, so it is best to stick with the far more widely deployed and popular time division multiplexing (TDM)–based technologies, such as Global System for Mobile Communications [GSM]).

Each of these can easily be addressed with counter-arguments. The first point is that while new services are always first marketed to early adopters at a premium, there is a natural spillover to a broader segment. Enterprise workers are accustomed to easily accessing information, such as e-mail messages and graphic intensive multimedia files, while they're in the office. Corporate information officers clearly see the productivity gains and benefits that extending these functions to their mobile workforce would have and are willing to pay for it. Proof that, if you build it, stay close to the market trends, and meet a critical need, customers will come to this field of dreams.

True, mobile high-speed data services are not yet available to mimic the on-line personal computer (PC) experience on mobile phones. But operators can sow the seeds *now* by exploring with device manufacturers new content and applications that take advantage of ever-faster transmission rates. Until there is ubiquitous coverage, they can provide "hot spot" high-speed data access with strategically deployed wireless local-area network (WLAN) technology. The trick is to match the device appropriately to the application and network capability and not to sell the 3G idea as an untethered PC experience.

Some new operators have abandoned the pursuit of 3G licenses or delayed buildout because they see no way of recouping their license and infrastructure investments. Others, thanks to regulatory relaxation, have devised network-sharing schemes or network hosting arrangements, commonly known as mobile virtual network operator

(MVNO) scenarios, to gain a fast entry. The U.S. market is a real-world example how new players have overcome the challenge of competition against incumbents through national branding and consolidation.

Lastly, the technical challenges faced in some quarters by 2G operators unfamiliar with spread-spectrum technologies can and will be overcome—sooner rather than later. To accelerate the return on investment (ROI) on 3G costs, operators need to focus on the immediate market opportunity: high-speed, mobile data services. That means weaning off current technology and getting to 3G *sooner* rather than later.

What's Different Today Is the Internet

The last decade has witnessed a phenomenon that offers the best counterpoint to the 3G naysayers: the advent of the Internet and the New Economy.

One might argue that the burst of the dot-com bubble weighs heavily against the prospects for the mobile Internet, but all signs point to the contrary. In fact, a global independent market study (August 2001) commissioned by Lucent Technologies demonstrated how the need for enterprise data communications is not being sufficiently met for the New Economy. To achieve cost-efficiencies, many enterprises are relying on an increasingly mobile workforce; yet, without the appropriate access to corporate data and communications resources, these workers feel increasingly isolated, frustrated, and out of touch with both their employer and customers.

The number of Internet subscribers worldwide is destined to continue its steady growth, reaching some 650 million by 2005, or nearly twice the number of 2000. Also, consider that over the next three years, the *average* global penetration of wireless communications will grow to 20 percent and that penetration will exceed 70 percent in key markets where wired Internet access is lagging. These trends are strong indicators that pent-up demand for always-on mobile Internet access won't remain just an enterprise-focused service for long (see *Figure 1*).

It is increasingly clear that 2001 was yet another of the historical cyclical technology transition years, when network service providers (NSPs) and the financial institutions that back them hedge their bets on next-generation services as they weigh multiple options for capital expenditures. The events of September 11, 2001, further complicated the decision process, because it clouded the prospects for economic recovery in 2002 and beyond.

The Murky Picture Is Now Clearer

The combined effect of a technology transition year in 2001 and the general economic slowdown, spawned by what's been referred to as "exuberant" market expectations gone sour, has created the current correction that is now manifesting itself in slower telecommunications capital expenditures (CAPEX). The global CAPEX pie for telecom is forecast to shrink by as much as $60 billion to $231 billion globally this year. To rid their core telecom business balance sheets of debt-laden mobile operations, some service providers have created separate, self-standing mobile enterprises or have spun them off completely. While committed to a still sizeable capital expense program ranging from $70–$80 billion globally, the mobile service providers are taking a far more sober approach to their business. In some cases, they have seized an opportunity to bulk up around converged standards to achieve economies of scale, yielding consolidation to the point where the top 20 mobile operators now control 80 percent of the CAPEX.

Most importantly, however, in light of the bursting of the dot-com bubble, the largest incumbent service providers have now changed their assumptions about what the future holds and must now plan accordingly. Unlike the two previous years, business decisions today can no longer be based on the assumption that equity is limitless and can be traded as cash to fund current operations. The large incumbent mobile operators have spotted the inevitable rise in Internet usage and are looking for *real-world* models for tapping into the underserved or unserved market segment for new services.

In approaching the opportunity, carriers are now very methodical. Specifically, their objective is to do the following:

1. ***Expand their most profitable customer base.*** By inference, this means putting a priority on business users and business applications that can be marketed to corporate management information system (MIS) decision-makers as a means for *increasing productivity* and *enhancing customer-relationship management (CRM).*

FIGURE 1

Mobile Internet Growth Prospects

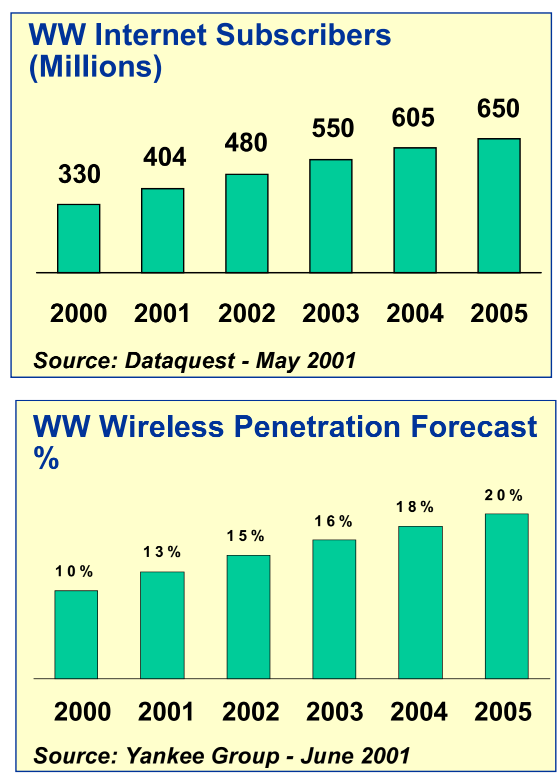

2. ***Create high-value, high-margin services.*** One-size-fits-all calling plans can only go so far in attracting customers, because with each additional customer, there is a diminishing rate of return. Carriers recognize that data-based services—including secure, virtual private mobile networks for extending enterprise communications and business functions—present the best opportunity for incremental revenues from enhanced services

3. ***Provide more capacity cost-effectively.*** The never-ending cycle of an expanded user base and need for more capacity to handle the increased traffic is a constant juggle for service providers. The latest intelligent antenna technology and network optimization tools offer the most cost-effective way to manage growth in capacity and coverage—without incurring the cost of new base-station cell sites and equipment.

4. ***Accelerate the implementation of mobile high-speed data technology.*** Operators and handset vendors that have marketed text-based short message services (SMSs) and clumsy text information retrieval on small mobile phone screens as a "mobile Internet" experience have left many early adopters less than satisfied. Rather than prolong this experience with interim solutions, carriers should look to *fast track the promise of mobile high-speed data*, provided vendors can deliver the solution cost-effectively.

5. ***Preserve embedded assets.*** With investment bankers, venture capitalists, and government regulators looking over their shoulders, mobile operators must opt for a solution that best meets their needs for maximum longevity, through multiple standards evolutions. Forklift upgrades can no longer be tolerated in today's tight credits market.

6. ***Turn in solid economic performance to drive up market capitalization.*** When all other criteria are combined, carriers need to execute their plans and create a business with intrinsic value. Higher valuation based on *real*, rather than *potential* performance, creates the capital needed to fuel subsequent incremental and measured growth and a solid economic foundation.

No doubt the challenges facing operators today are formidable, but with the right technology partner, they are surmountable. The key is to develop a market-driven, cost-sensitive business plan and implement a service intelligent architecture, which takes the latest hardware and software for the end-to-end network and makes sure that advanced services can be easily provisioned, maintained, and billed through a highly integrated software management platform.

Key Considerations for Success

Where does the operator start on the road to success? The best point of departure to a solid business scenario is the selection of a strategic partner who understands and can best help execute the following:

1. ***Fast-track to mobile high-speed data services.*** The world has converged on spread-spectrum radios as the common technology for 3G. By extension, the lowest risk path for any operator is to partner with someone with the most extensive global experience in designing and building commercially successful spread-spectrum–based network infrastructure. The most proven and commercially successful spread-spectrum networks today are based on CDMA.

2. ***Maximize the power of the network.*** Unlike other access providers whose roles are relegated to a "dumb pipe," mobile operators have the advantage of seizing the power of network intelligence to extract value that can be packaged for premium, value-added services. A systems integrator well versed in service provisioning can help any operator leverage network intelligence to extract value from this asset.

3. ***Pick the lowest-hanging revenue fruit.*** A Yankee Group study published in December 2001 quantifies the Lucent-commissioned research, which indicated the pent-up demand for extending corporate intranet and e-mail communications seamlessly to mobile workers. The Yankee Group predicts that in the United States alone there will be about 43 million wireless enterprise data users by 2005. It is safe to say that the number of potential users worldwide will be even greater. That's a ready-made, real-world market and a huge revenue-generating opportunity for any service provider.

4. ***Turnkey, out-of-the-box solutions.*** In addition to tightly integrated and easily managed network elements, network operators need to rely on strategic suppliers that have ready-made, out-of-the-box solutions that they can market to the premium enterprise customer. This encompasses a range of standards-based enterprise business solutions that support corporate functions and processes. In the New Economy, businesses must enable mobile workers to be as productive and responsive to customers as if tethered to their office resources. Additionally, deploying these capabilities will create the environment in which services for other market segments can arise and flourish. Moreover, providing always-on, always-accessible CRM systems are critical to remaining competitive in the New Economy. The right strategic partner will have packaged solutions ready for deployment—from back-office support systems, applications software support, and compatible wireless handheld devices optimized for the appropriate function.

5. ***Speed to market.*** The faster the deployment, the faster revenues are realized to shrink the ROI interval for the operator's CAPEX. This calls for rigorous professional business and engineering services that encompass all planning aspects through deployment, launch, and beyond. By relying on this disciplined support, the carrier can build a credible, real-world business plan that establishes accountabilities and financial performance levels that are sure to lead to a successful track record to warrant higher valuations and economic health.

Conclusion

The sky did not fall down on the mobile sector in 2001. It was, in retrospect, part of the normal cyclical downturns that are typically initiated by generational changes in technology. Unfortunately, this trend was exacerbated by the burst of the dot-com bubble in the credit markets and the economic uncertainty following the events of September 11. However, by focusing on business fundamentals and a disciplined methodology for value creation, mobile operators can indeed anticipate sunny skies ahead. The right strategic relationship with a reliable and experienced partner can make all the difference.

The Impact of Intelligent "Push" Services on Wireless Carriers

Jeff Condon

Vice President, Marketing

Otelnet

I. Executive Summary

The world of mobile phones is changing rapidly and dramatically, with global utilization approaching a billion subscribers. Until recently, mobile phones' primary use was for "phone calls"—person-to-person voice communications. With the unprecedented success of person-to-person short messaging service (SMS) in Europe and the iMode phenomena in Japan, the use of the mobile phone is shifting from a rather basic voice communications tool to a true "mobile terminal" whereby the user may have access to the plethora of data otherwise accessible only by those with personal computers (PCs) connected to the Internet. The thirst for information, as well as other data-oriented services, while mobile, has not been quenched by the use of the micro browser—"pull only" services—on the mobile phone. The market has learned that bringing Internet services directly from the PC to the mobile phone has not worked, but valuable information and services can be brought to the mobile phone. Pull services alone are not sufficient. This has created a major opportunity for wireless carriers to better serve their subscribers and to gain a significant new source of revenue by intelligently pushing desirable, valuable information and services to their subscribers.

II. Wireless Data, Content, and the Mobile Subscriber

A. Success of SMS in Europe As a Precursor to Push Services

The success of SMS in Europe, a bit of a surprise since it grew "organically" with little or no marketing of the service by the wireless network operators, has set the stage for consumer uptake of new information or content-oriented services. The ease with which subscribers, most notably younger subscribers, send many millions of short text messages each month to one another has created a very receptive audience for receiving text messages from other sources. Today, something on the order of 90 percent of all SMS messages sent and received in Europe are person-to-person, or phone-to-phone. This is due in part to the ability of conversing with another person while avoiding the ringing phone and by avoiding the cost of making a phone call.

The net result is a huge consumer market already accustomed to receiving (and sending) text messages. A natural progression, then, is to "feed" this hungry market with more of what it wants. Thus, sports scores, for example, are readily available to be sent (pushed) to enthusiasts, often within minutes of a score change, for example, in a football game. Standard content services, such as these sports scores, daily horoscopes, generic news items, and local weather reports, have become "checkbox" items. Since these phones, and the supporting networks, allow the consumer to send messages as well as receive them, new services are becoming available to take advantage of this well-primed market: interactive voting, lifestyle and entertainment services such as gambling, purchasing and downloading of ring tones that have been created from popular music, and so on are leading the charge. This hungry market is willing to pay a bit more to get interesting information, or content, pushed to them if the content is perceived to have high interest or high value.

SMS is an estimated $1 billion a month industry in Europe, although it has lagged significantly elsewhere due to the lack of inter-carrier solutions. In the United Kingdom, for example, SMS message traffic increased more than 700 percent in the nine months after an inter-carrier solution was deployed, according to the GSM Association.[1]

As shown in *Figure 1*, SMS traffic in Western Europe shows just how hungry this market is. For example, in 2001 the average mobile phone user sent or received about 64 individual short text messages each month, and that number is projected to continue to grow through 2005 (upper line). Meanwhile, the number of SMS messages per mobile phone indicates that these users actually have more than one mobile!

This data is based on a mobile phone penetration rate of about 70 percent in the top 15 European counties, meaning that of every 100 people in these countries, 70 of them use a mobile phone.

B. How the European Market Is Different from North America

There are some key differences between the wireless phone markets of Europe and of North America. These differences account for the strong uptake of SMS and SMS–based information services in Europe, whereas this is just beginning in North America.

FIGURE 1

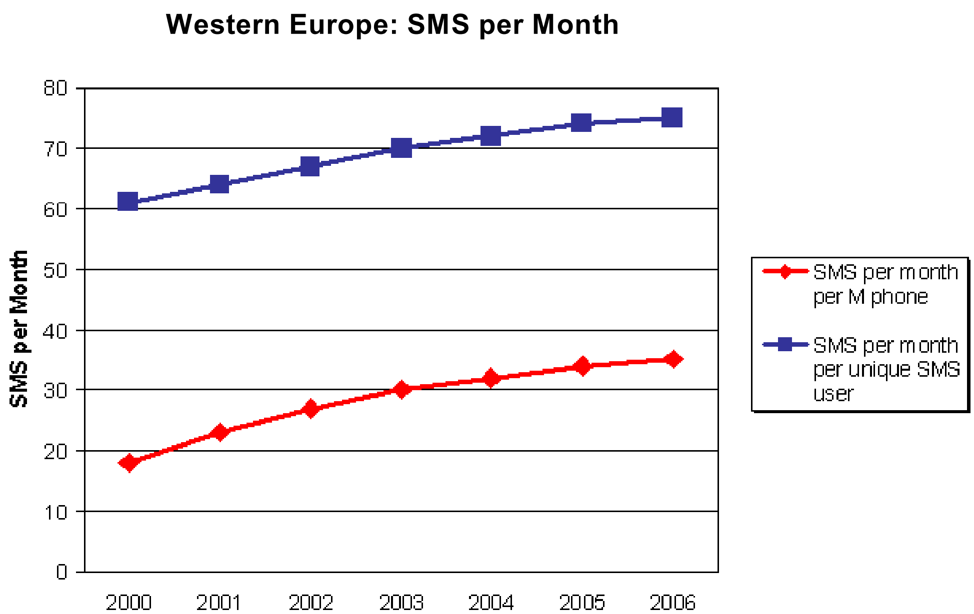

Source: Research Centre of Bornholm (Denmark) + OTELNET estimates

European wireless networks are nearly all Global System for Mobile Communications (GSM), due to a conscious decision by network operators and governments in the mid-1980s to implement a common network protocol for digital (second generation, or 2G) mobile phone networks. One reason this decision was made and agreed to was to make it easy to add new, high-value services, as well as to facilitate the migration to general packet radio service (GPRS) (2.5G) and to Universal Mobile Telecommunications System (UMTS) (3G) packet-switched networks. This resulted in 2G networks being implemented throughout Western, Northern, and Central Europe based on the GSM standard. Having multiple network operators (usually three or four) in each country building GSM networks greatly facilitated voice call roaming, since all that is needed are service and billing agreements among the operators. Thus, a subscriber from Helsinki can travel to virtually any European city and be assured of service availability. This holds true with SMS as well. One can easily send text messages to a subscriber in another country without worrying about incompatible networks. This ease of communications throughout Europe has had a major influence on the high penetration rate of mobile phones in these countries.

In North America, by contrast, the migration from first-generation analog mobile phone networks to 2G occurred in a much different way, and much more slowly. The philosophy shared by European operators to build homogeneous networks to ensure good service and easy roaming was not used in North America, where competition drove the technology choice, not government regulation; North American network operators built 2G networks in the least expensive way they could. The result was a much slower build out of 2G digital services and heterogeneous networks. These networks were based mostly on time division multiple access (TDMA) and code division multiple access (CDMA), with only a few operators choosing to build a GSM network. As a result, when one subscribes to a digital mobile phone service in North America, the handset is capable of working only on that type of network—CDMA for example (these handsets usually can also be used on 1G analog, or advanced mobile phone service [AMPS], systems)—due to the length of time it has taken to build out the digital networks. This has had a dramatic effect on roaming, since a CDMA handset can only roam on a CDMA (or AMPS) network. Of course, these heterogeneous networks have made it difficult to implement cross-network SMS services.

The difference in network evolution has lead to a much slower uptake of mobile services in North America. Key factors driving mobile phone penetration are network homogeneity, life styles, service availability, and price. Only now is the U.S. market approaching 50 percent penetration, meaning that about half the population has a mobile phone. Interestingly, the percentage of American homes that have a PC with an Internet connection has passed 70 percent. Thus, a large majority of American consumers have easy access to information—content—via a "fixed" Internet connection. European countries, by contrast, have a much higher penetration of mobile phone users, and they are all using the same type of device (GSM). This means that mobile network operators, or carriers, in North America still have half the population to sell basic services to and have strongly concentrated new revenue-generation activities on acquiring new subscribers. European network operators, in contrast, do not have this "green field" of population to sell to—most of the available European market (those older than about 9 years of age!) already have a mobile phone, so European operators must find other ways to generate new revenue.

C. The Wild Card: Japan

NTT DoCoMo's success with its iMode service is an interesting and unique phenomenon. Although NTT is not the only wireless provider in Japan, it is the best known.

There are interesting circumstances surrounding iMode and the Japanese market. The service offerings include both push and pull services, are graphical in nature, and are "always on." In addition, the mobile phone penetration is approaching 65 precent; there is a low PC penetration rate; the culture is strongly entertainment-oriented; and the DoCoMo business model of revenue sharing provides strong incentives for content and applications providers. The Japanese consumer desires entertainment-oriented services, often taking the form of small, animated graphic characters. This same consumer most likely does not have a PC in the home and very limited or no Internet access.

The iMode service fulfills the demand for content that is currently satisfied by Internet access via the PC in the Western cultures of Europe and North America. In addition, NTT has created a business model that provides strong incentives for third-party application and content providers to deliver offerings through the DoCoMo iMode service: A significant majority of the fees paid to DoCoMo by the consumer (more than 90 percent) are passed along to the third parties. The result is a plethora of high-quality service offerings for the Japanese mobile consumer to choose from. This has created an interesting competitive situation with other Japanese mobile network operators whose network and service offerings are more like those in Europe.

DoCoMo now has investments in both Europe and North America, and both markets will learn and benefit from the DoCoMo experience. Since lifestyle, cultural, and service-availability factors are very different for Japan, it remains to be seen how well the iMode model will be accepted outside of Japan.

III. Why "Pull" Services Have, So Far, Not Done Well

A. What Is a "Pull" Service?

"Pull" services, in this context, refer primarily to the paradigm of implementation of the "fixed" PC, with a browser and wired Internet connection. The consumer (subscriber) uses the browser and other "on-line" tools, such as search engines and portals, to seek out and "pull" content to the PC.

The first attempts to make Internet content available to mobile "terminals" (e.g., mobile phones and, to a lesser degree, personal digital assistants [PDAs] such as the Palm VII) was, and still is, intended to implement this same paradigm on the mobile phone through a browser—actually, a "micro-browser" on the handset, with which one attempts to "surf the Web" as if it were a PC–based browser.

B. WAP and iMode

Wireless application protocol (WAP) and iMode are two competing network "transport" technologies that provide a mechanism for data to move between the mobile phone and the network. WAP has been most widely implemented in the West (Europe, North America); iMode is the technology used by NTT DoCoMo in Japan, as was discussed. There are some key differences between the two technologies. The current implementation of WAP is on 2G, circuit-switched networks. This requires the user to make a phone call (i.e., dial a phone number) to create the "circuit," or connection, to the network before the micro -browser, often called a "WAP browser," can be used to search for, or "pull," information and content. These micro-browsers are text-only implementations and are forced to communicate over the very slow data circuits of a 2G network. Furthermore, the user pays the same per-minute charges as for a voice call.

iMode, by comparison, is implemented on a packet-switched, "always-on" network. Although slow by broadband Internet standards, iMode is faster than circuit-switched WAP and does not require the user to make a phone call to set up the connection to the network. The iMode browser has been optimized for its network, so the user experience is more satisfactory than for WAP implementations. Billing is different as well. Where the WAP user currently is billed per minute of connect time, the iMode user is billed per data packet received or sent by the handset, regardless of connect time. In fact, the concept of connect time does not apply, since the handset is "connected" to the network whenever it is turned on, much like a PC connected to the Internet via a digital subscriber line (DSL) or cable modem.

Keeping in mind the differences in cultural drivers between Japan and the West, plus the organic success of SMS in Europe, one can easily see why micro-browser–based, pull–based services have failed in Western markets and that iMode services have seen a resounding success in Japan. It is the user experience that will determine success or failure, and the user experience with WAP–based browsing has been very poor. A combination of factors has caused this, including that the user interface—a text browser on the very small screen of a mobile phone—is awkward and difficult to use. The very slow connection to the network significantly exacerbates this awkwardness, since each time the user sends, for example, a menu choice for a new screen of information, the latency, or waiting time, for the new information is agonizingly slow. For a time, wireless network operators tried to mitigate this by creating the "walled garden," which essentially restricts the user to the browsing menu choices provided by the operator. Unfortunately, this worsened the situation by preventing the user from getting at information or content that is readily available on the Internet-connected PC, so this approach has mostly been abandoned. The result is that most mobile subscribers who have tried to use WAP and micro-browsers rapidly abandon it in frustration.

In Japan, the iMode browser is more intuitive to use, and the mixed graphics and text appeal to the Japanese consumer. Also, the typical Japanese mobile subscriber does not have ready access to the Internet via a PC, and so has no preconceived notions about the user experience of browsing.

IV. The Evolution of SMS

Infrastructure refers to the mobile network and how the mobile phone user actually can use the network for something other than making phone calls. This is not an attempt at a dissertation on 2G, 2.5G, and 3G networks, but some discussion is needed to understand why "push" services are so important to wireless network operators and subscribers alike.

The majority of wireless networks have been migrated to digital service using the three aforementioned technologies: GSM, CDMA, and TDMA. Each of these has unique characteristics, with advantages and drawbacks, and it appears that TDMA systems are nearing the end of their life, such that any new 2G networks are either GSM or CDMA. In any event, all three have the common characteristic of being circuit-switched, meaning that for the user to connect to the network, a phone call must be made to create the connection, or circuit, to the network. This is fine for the most part for making simple voice calls but is not fine for data.

As mentioned, this connection (regardless of whether its GSM, CDMA, or TDMA) is so slow as to have very limited usability. One reason that SMS has become so successful in Europe is that it does not use the "circuit" and does not require a phone number to be dialed. Rather, is uses a different part of the network, called signaling sstem 7, commonly called SS7. This is the part of the network that, among other things, causes your phone to ring when someone places a call to you. This is true for today's wireline phone networks as well as wireless networks. This process of making your handset ring is accomplished before creating the circuit from your handset to the network and to the calling phone.

Without this SS7 mechanism, the process of initiating a call—dialing a number and making a handset ring—would completely fill and overwhelm today's networks. SMS uses the SS7 network to send short messages to (and from) the mobile handset. As it happens, most network operators have also implemented e-mail gateways between their wireless networks and the Internet; the result is that one may send an e-mail to most (SMS–capable) mobile phones, and some phones can originate an e-mail message as well.

As described, about 90 percent of the SMS traffic in Europe is handset to handset—all done without making a single phone call and without using up voice call minutes, and all via homogeneous GSM networks that have roaming agreements not only for voice calls, but also for SMS traffic. This is one of the reasons why SMS has been so successful and has created a major opportunity for wireless network operators to create new revenue-generating services that leverage this capability.

V. Content Sources, Wireless Portals, and the Value of Intelligent "Push" Services

This opportunity for wireless network operators revolves around creating new content-oriented services. If the operators can capitalize on this now, they become well positioned to take advantage of the increased bandwidth (with more than just greater voice call capacity) to be made available with the migration to packet networks: 2.5 G and 3G. Pull will get better, and push will still be needed for many services. This means getting access to sources of content—content that the network operators cannot create on their own. Virtually all wireless network operators now offer some sort of portal, mostly accessed via PC–based web browsers, and offer mobile phone micro-browser access to some of the portal content, which has been specially formatted and often replicated for the micro-browser. The larger multinational operators have created separate lines of business in support of their portals, using this increased buying power to create relationships with a variety of content partners, including content aggregators and large original content sources such as international news services.

Smaller operators have outsourced their portal to wireless application service providers (WASPs). This content is made available to the operators' subscribers, either branded by the content provider or by the operator's portal. This model is very similar to, for example, the Yahoo! portal. Also like Yahoo!, the operator's customers can subscribe to have a limited amount of content "pushed" to their wireless phones as SMS messages. For the most part, the kinds of content currently offered to be pushed to the mobile handset are regularly scheduled standard content types, such as the aforementioned horoscopes, general weather conditions, and news "snapshots." This kind of content is useful but has not taken on the cachet of "valuable" to the extent that the wireless consumer will pay a premium for it.

Intelligent push services become important to the subscriber for key reasons: convenience, more personalization than via a browser, and event-driven delivery. As wireless subscribers, especially in Europe, become accustomed to these kinds of content services, usually on a subscription basis either at no extra charge or for a small monthly fee, and receive information on a regular basis rather than having to poke around with a wireless browser to find it, a significant revenue opportunity comes into play: Offer content, or information, often on an event or "triggered" basis, for which subscribers are willing to pay a premium price because the information is perceived to be valuable.

VI. The Nascent Opportunity for High-Value, Premium "Push" Services

With the uptake of SMS–based push services beginning, following close on the heels of the person-to-person SMS phenomena in Europe, wireless carriers—and other players in the wireless services value chain—are beginning to explore "premium services" or "premium alerts."

Since, as described, the European (and Japanese) markets provide a fertile, receptive audience for information to be pushed to the mobile device, "premium" alerts are relevant, personalized, and perhaps location-sensitive–compared to "standard" alerts, which are more general and generic in nature. *Table 1* shows some examples of each.

A. What Makes a True Premium "Push" Service Offering?
Intelligence for push services means being relevant, personalized, and timely. If an alert (a message, some information, some content, etc.) is to be pushed to a subscriber such that the subscriber is willing to pay extra—perhaps a lot extra—for it, then the decision to send the alert must be based on a number of attributes of the subscriber and the content. The subscriber should not have to engage in a tedious search process for information; rather, a set of criteria needs to be pre-established to make it extremely easy to use. This is a key reason why "push" is such an important content-delivery model. The service provider must provide an easy to use interface for the user to create subscriptions to all sorts of alerts (self-provisioning), as well as to create and modify a profile of characteristics that is used to decide when a given alert is to be sent (pushed) to the mobile phone.

TABLE 1

"Standard" Alerts	*"Premium" Alerts*
Sports scores	Real-time, event-driven alerts
Stock-market data	Opt-in advertising,
Ski reports	M-Coupons
Weather	Location- and presence-sensitive alerts
Where the fish are biting	Actionable alerts (e.g., making a purchase)
Headline news	One-2-one marketing
Horoscopes	Profile-based services

Subscriber attributes: For "scheduled," standard alerts, this is fairly easy to implement and has been done so by most network operators, usually within their portal, and often outsourced to a WASP. Significantly more difficult is being able to offer event-based alerts that must be sent only when all the elements of the user's profile have been satisfied: time of day, day of week, time zone, presence on the network (i.e., whether the mobile phone actually turned on), physical location or proximity, and others.

Content attributes: The alert must be *intelligently pushed* based on all of the scheduled alert criteria that the user chooses plus the variables of event-based content (a stock hits a high or low, and the user needs to react quickly; a traffic crash just happened one kilometer ahead, and the next highway exit is 500 meters ahead; a poisonous gas just escaped from the chemical factory, and the cloud is moving to the user's position; the user has just passed near a Macy's and is sent a coupon she "opted in" for to get a special price on a new pair of jeans, but only on Thursday evenings, etc.). Clearly, the user cannot be interactively actively searching (pulling) for these events; these kinds of services are highly valuable to the user when *intelligently pushed* to their mobile phones under just the right conditions.

B. What Does the Carrier Need to Effectively Offer Premium Services?

Although the potential for significant new revenue is great, intelligently pushed content, messages, and alerts put a lot of requirements on the wireless carrier. The necessary software is complex and sophisticated: It needs to constantly perform matchmaking between myriad content sources and millions of subscribers, each with a set of very specific attributes. A robust and flexible rules engine would allow subscribers to specify the information delivered to them such that it is highly personalized and timely and delivered only under the conditions dictated by the subscriber. They also need access to standard, easy-to-use Web interfaces to set up and modify their user profile, notification services, and delivery preferences.

An additional requirement is that the software environment must make it easy for the carrier to implement new content services quickly and to easily add (or remove) and manage new content sources and types. The software must be able to track and account for every message sent and update the carrier's billing systems. It needs to be easily scalable as more users and more services are added. And, it must easily integrate with the systems in which the carrier has already invested. In short, this must be "carrier-grade" software. It is unlikely the carrier has the resources to build this in house (beyond perhaps building a prototype) and will need one or more vendors and systems integrators to do the job properly.

VII. How Carriers Can Make Significant New Revenues—and Why They Must

With all of this background, what does the revenue opportunity look like for the carriers and for other participants in this value chain? Is it significant enough to justify a substantial investment in this intelligent "push" technology and to create the partnerships needed to make it work? Are the carriers satisfied with their current means of generating new revenue?

As described, existing opportunities for new revenue are different in each of the major geographies. Dramatic uptake of intelligently pushed services is occurring new in Europe and will continue to grow. Since the European wireless market is becoming saturated, carriers serving this market are desperately seeking new and innovative revenue sources, since they can no longer depend so heavily on gaining new subscribers. Obviously, they have experienced the SMS phenomenon and want to somehow capitalize on it. They tried the "Web browsing" paradigm, using microbrowsers and WAP, and it has not yet worked well. As described, the next logical step is to ride the SMS wave by offering creative new services delivered via SMS. This is particularly attractive as the migration to packet networks (GPRS) is under way. Both the network operators and the wireless portals are now developing and launching new services, many focused on entertainment as well as affinity and permission-based "m-marketing," with premium price tags—and subscribers are biting.

The Japanese market, with its successful iMode service model, continues to grow at a rapid pace, and new push-and-pull services are rapidly being introduced that take further advantage of this receptive market and the technology base.

The North American market, especially the United States, will have the unique advantage of learning from both

Japan and Europe and applying this learning to the rapidly growing youth market that has migrated from carrying pagers to carrying mobile phones and are rapidly embracing person-to-person "texting." Similar to the European SMS phenomenon, this is creating a large receptive market for push services.

VIII. Conclusions

The SMS phenomenon in Europe, the success of iMode in Japan, and the convergence on 2.5G and 3G networks in all markets have created a hungry and receptive market for content, or alerts, to be intelligently pushed to mobile subscribers' handsets. To capitalize on this huge opportunity for badly needed new revenues, the carriers must take clear action to implement technology and operational practices that will provide them with the means to deliver valuable, desirable content to their subscribers in a way that takes advantage of existing investments in networks, software, and human resources.

A key element is to find vendors the sophisticated software and services that are needed and to implement them in a way that maximizes revenue to the carrier. In most cases, this means the carrier must implement in the carrier network rather than outsourcing and giving up a large portion of the new revenue to the outsourced service provider. Fortunately, this new genre of carrier-grade software is beginning to appear on the market. One important point is that the user experience for pull services will likely improve as higher-bandwidth packet networks—2.5G and 3G—become available and as pull services created uniquely for the small "footprint" of the mobile handset are brought to market. The most likely forward-looking scenarios are services that are implemented according to how they work best: some pull, many push, and many a mix of both.

Notes

1. Source: INT Media Group

Wireless Data in the Coming Decade

Jay English
Consultant, Network and RF Systems Engineering
inCode Telecom

Introduction

"Is that all there is?" Astronaut Al Bean's musing to Mission Commander Pete Conrad upon departing the moon during the Apollo 12 mission seems to fit the mood surrounding wireless data in the latter part of the 20th century. For all the tremendous technological advances the industry and the world saw during the previous decade, we were left feeling like there should have been something more. The promise of 3G with "always-on" high-speed wireless connections, fixed wireless broadband to home and office, and wireless local-area networks (LANs) removing the leash of cable from our laptops and personal computers (PCs) led to high expectations and an almost palpable anticipation. In spite of all the promise, excitement, and investment made, the century turned without those services making an appearance on a wide scale, and, despite tremendous technological advances and subscriber growth, so did the market.

2001 saw the beginning of man's permanent habitation of space, a concept that was only a dream when Bean and Conrad stepped foot on the moon a scant 30 years ago. It also saw what could only be described as one of the most horrific events in human history on September 11. No small wonder that somewhere in between the two our economy tumbled with technology companies, and telecommunications in particular, being hardest hit.

While the growth and promise of high-tech opportunities teetered on the brink of disaster, the determination and resolve of the wireless data industry refused to be quelled. In spite of the odds, lack of capital, and a sagging economy, bold steps were planned to forge ahead. 3G1xRTT was launched in overseas markets, wideband code division multiple access (WCDMA) and Universal Mobile Telecommunications System (UMTS) groups continued development plans, U.S. carriers announced their intent to implement a variety of data services in the coming decade, and market forecasts showed data overtaking voice and consuming an ever-increasing portion of the wireless market by 2010.

This paper is an examination of the possibilities and applications that await us in the coming decade in the world of wireless data. We will examine 2.5 and 3G services; discuss the differences between high-speed circuit-switched data (HSCSD), short message service (SMS), and packet data; and study the offerings available. Those offerings range from relatively inexpensive overlays like cellular digital packet data (CDPD) to next-generation services such as WCDMA and UMTS and include a few other increasingly popular options like SMS, enhanced data rates for GSM evolution (EDGE), and general packet radio service (GPRS). We will discuss what this means for carriers and consumers and make a few "educated guesses" as to what the wireless landscape might look like as we move forward into this first decade of the 21st century.

Current industry predictions show a strong trend toward dynamic growth in the wireless data sector. Currently, wireless data occupies approximately 2 percent of the marketplace. However, it is projected that more than 60 percent of all wireless subscribers will utilize some form of wireless data by 2007 (see *Figure 1*).

The technologies that will be implemented to enable this growth include current HSCSD, CDPD, GPRS, EDGE, and third-generation (3G) services such as 1xRTT/3xRTT, WCDMA, and CDMA2000. Each of these technologies offers unique services and benefits to both users and carriers and will be discussed in the course of this paper.

High-Speed Circuit-Switched Data

Initial efforts to provide data services to the consumer included HSCSD, which utilizes existing carrier infrastructure, with minimal modification and equipment addition, to provide the customer data connectivity over the traditional circuit-switched path. The addition of an interworking function (IWF) to the circuit-switched path allows for validation and routing of dedicated data traffic such as fax, file transfer, and Internet services over existing wireless networks. The maximum data-rate predictions for this service are 14.4 kilobits per second (kbps). Compression technology allows the user to experience speeds equivalent to approximately 28.8 kbps. Implementation of working networks was accomplished successfully by all major U.S. carriers by late 1999/early 2000. HSCSD call flow mirrors a traditional wireless call in its origination from the wireless device, connection over the air to a base station, transition to landline, and connection to the base-station controller (BSC) and mobile switching center (MSC). Where a traditional voice call would continue via the MSC to the public switched telephone network (PSTN), the data call will route to the IWF for further processing and routing to the end user or application. End users can include fax machines connected via

FIGURE 1

Wireless Data Projected Growth Relative to Voice

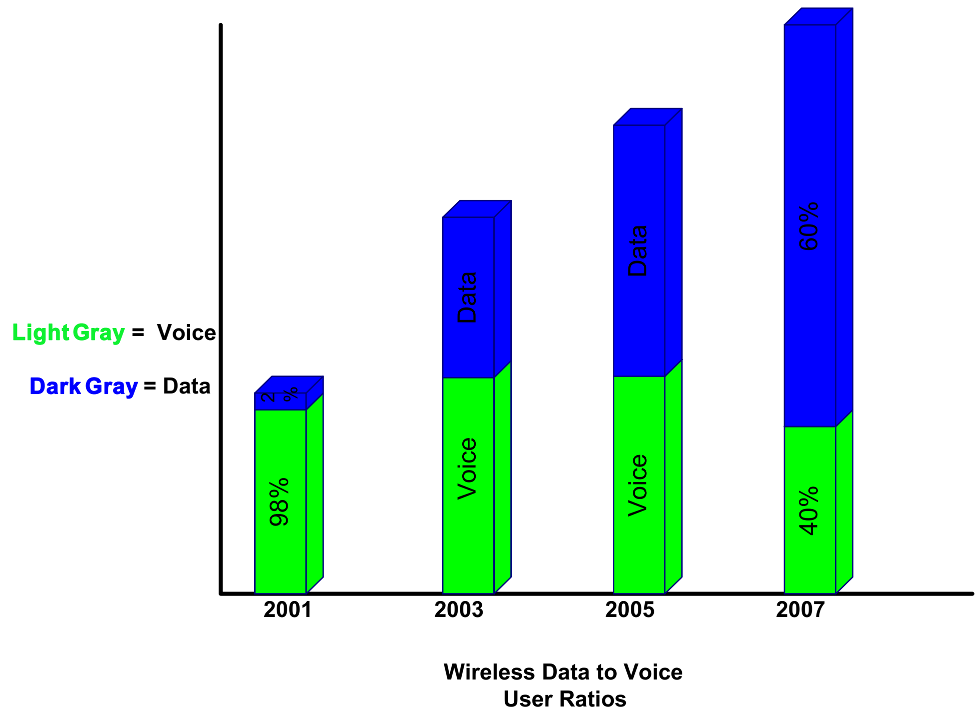

the PSTN, file transfers, and "Web surfing" via Internet connectivity along with virtual private network (VPN) connections via internal and external routers. An example of current HSCSD network configuration is illustrated in *Figure 2*.

The advantages offered by this service are the ability of the consumer to access services previously unavailable over wireless networks, such as e-mail and Web access. The major drawbacks include the impact on the total capacity of the carriers network, the low speed limitations, a result of the air interface, and the fact that users are charged for minutes of use as they must connect, conduct their transactions, and disconnect from the network, just as they do with voice calls. These limitations were noted from the onset of HSCSD, and new technologies evolved rapidly to overcome them.

Short Message Service

In addition to HSCSD services, carriers implemented another effective data transfer mechanism during virtually the same timeframe as HSCSD. SMS is a wireless service that enables the transmission of text-only messages between mobile subscribers and external systems, such as electronic mail, paging, and voice-mail systems. SMS provides a mechanism (via signaling system 7 [SS7]) for transmitting short messages to and from wireless handsets. The service makes use of a SMS center (SMSC), which acts as a store-and-forward system for short messages. The wireless network provides for the transport of short messages between the SMSCs and wireless handsets. In contrast to existing text message transmission services, such as alphanumeric paging, the service elements are designed to provide guaranteed delivery of text messages to the destination. A distinguishing characteristic of the service is that an active mobile handset is able to receive or submit a short message at any time, independent of whether or not a voice or data call is in progress. SMS is characterized by out-of-band packet delivery and low-bandwidth message transfer.

Initial applications of SMS focused on eliminating alphanumeric pagers by permitting two-way general-purpose messaging and notification services. As technology and networks matured, a variety of services were introduced, including electronic mail and fax integration, paging integration, interactive banking, and information services such as stock quotes. Wireless data applications include downloading of subscriber identity module (SIM) cards for activation, debit, and profile-editing purposes.

The benefits of SMS to the service provider include the following:

- Increased call completion on wireless and wireline networks by leveraging the notification capabilities of SMS
- An alternative to alphanumeric paging services
- Enabling wireless data access for corporate users
- Provisioning of value-added services such as e-mail, voice-mail, and fax/mail integration; reminder service; stock and currency quotes; and airline schedules
- Provisioning of key administrative services such as advice of charge, over the air downloading, and service provisioning.

SMS comprises two basic point-to-point services: mobile originate (MO) and mobile terminate (MT). Mobile-origi-

FIGURE 2

Example of Typical HSCSD Network

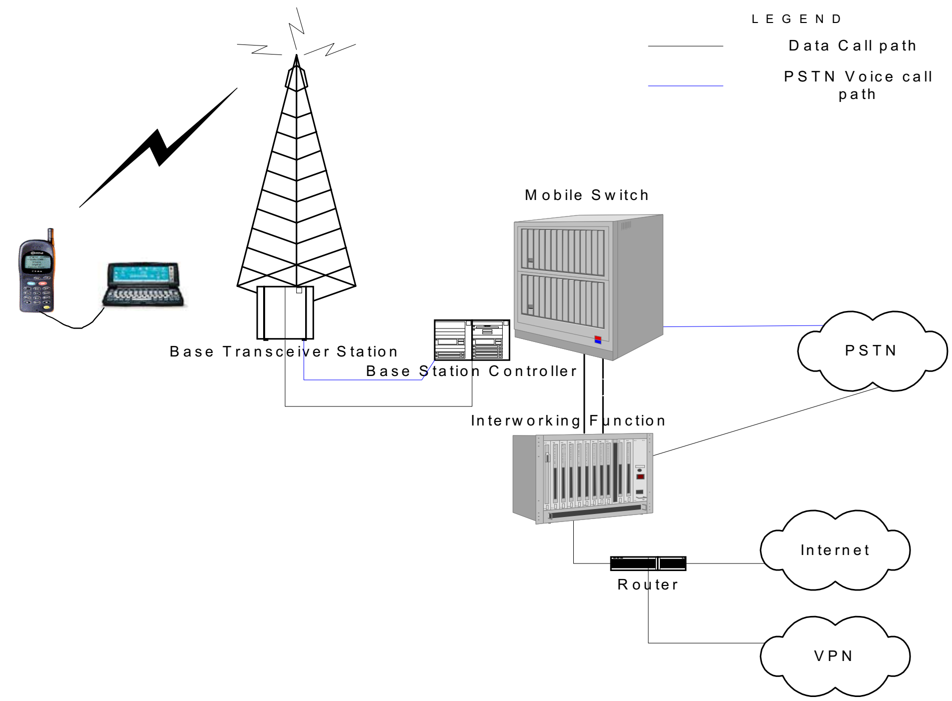

nated short messages (MO–SMs) are transported from the handset to the SMSC and can be destined to other mobile subscribers or for subscribers on fixed networks such as paging networks or electronic mail networks. Mobile-terminated short messages (MT–SMs) are transported from the SMSC to the handset and can be submitted to the SMSC by other mobile subscribers via MO–SM or other sources such as voice-mail systems, paging networks, or operators.
Depending on the access method and the encoding of the bearer data, the point-to-point SMS conveys up to 190 characters. Many service applications can be implemented by combining these service elements. Besides notification services, SMS can be used in one-way or interactive services providing wireless access to any type of information anywhere. Leveraging emerging technologies that combine browsers, servers, and new markup languages designed for mobile terminals, SMS can enable wireless devices to securely access and send information from the Internet or intranets quickly and cost-efficiently.

Intercarrier interoperability is often considered the "holy grail" of SMS. In Europe, where all carriers operate the same technology, interoperability has not been a problem. Usage of SMS in Europe has seen tremendous growth, especially within the teenage market.

In the United States, due to the dissimilarity of technologies in use, interoperability for SMS traffic has not been achieved. Current solutions, not yet deployed, include a "clearinghouse" approach in which all carriers "subscribe" to the clearinghouse routing services and "interoperability gateway" solutions. Gateway solutions convert the SS7 traffic to an Internet protocol (IP) backbone and transfer to the partner networks gateway, which then converts traffic back to SS7–compatible traffic for transfer back to a receiving SMSC. Numerous problems exist with both solutions, mostly that the carriers in the United States will not agree to a common solution.

In addition to the interoperability issue, compatibility with next-generation services is essential. GPRS and EDGE compatibility is a "*must*," as two major carriers in the Unites States—AT&T and Cingular—will be completing Global System for Mobile Communications (GSM) overlays to their time division multiple access (TDMA) networks. The intent of both carriers is to evolve to packet data solutions, either GPRS or EDGE, in the near future. Both carriers will seek SMS compatibility with their next-generation networks.

Plans for 3G1xRTT, 3xRTT, CDMA2000, and UMTS (WCDMA) also include compatibility for SMS traffic over the packet data architecture.

Cellular Digital Packet Data

The next evolution of wireless data technology is the transition from circuit-switched to packet-switched data services. In a packet-switched network, a data stream is broken into packets that are wrapped with information such as destination, length, and order (usually specified in a packet header). Each packet is transmitted separately and can take different routes from the sender to the recipient. The recipient receives these packets and uses packet header information to verify successful packet transmission and reassemble the packets into the original data. Packet-switched networks enable billing schemes based on the amount of data transmitted, not minutes used. Among the early entrants to packet data is CDPD.

The CDPD concept was originally conceived by IBM researchers in the early 1990s. In 1992, leading cellular carriers formed a standards group to create CDPD specifications, and shortly thereafter, in 1995, initial deployments of CDPD in North America began. Service availability now exists in the United States, and CDPD deployments also exist internationally in markets such as Mexico, Canada, New Zealand, Israel, and Venezuela.

CDPD is an overlay to existing advanced mobile phone service (AMPS)/TDMA networks and again builds on the existing infrastructure with minor additions and modifications. CDPD is still dependant on the traditional network for radio frequency (RF) coverage and is extremely vulnerable to frequency planning considerations and interference. The initial design of CDPD involved the creation of an overlay network consisting of mobile database stations (MDBS) and intermediate systems (MD–IS) to facilitate the packetization of data, and the utilization of unused voice spectrum. CDPD is designed to take advantage of the fact that much of the spectrum in existing voice networks is unused for a significant portion of time. CDPD utilizes the available spectrum on a per-timeslot basis to send short bursts of packetized data over the network. CDPD utilizes this “frequency hopping” scheme to scan for available slots on each channel and allocates packets for transmission based on the results of the scan. Early CDPD networks relied exclusively on this channel hopping to allocate and transmit/receive traffic. Early results indicated that packet switching became very limited in its usefulness as the capacity on the network increased. As a result, many CDPD carriers chose to utilize a dedicated CDPD channel. In this implementation, CDPD is allocated to a dedicated channel, and multiple users can be placed on the channel. The division of the packets and ability to avoid collisions is based on the complex Reed-Solomon calculations. CDPD’s reverse link is a contention-based system utilizing digital sense multiple access with collision detection (DSMA–CD) to determine if the allocated channel is busy. Two possible events can occur:

1. *Idle Channel:* If an idle channel is encountered, then the CDPD modem immediately begins to transmit data packets. The modem then waits for an acknowledgement from the base station:
   - *Successful Acknowledgement:* Packet has been received correctly by the base station
   - *Decode Failure:* Packet was received by base station but was not decoded correctly (usually because another modem begins transmission at the same time). The CDPD modem assumes a collision has occurred and “backs off.” Transmission is retried after a random delay time.
2. *Busy Channel:* If reverse link is busy, then the CDPD modem “backs off” and delays transmission for a random amount of time.

If a collision is detected, the CDPD specification uses an exponential weighting system to determine when a modem should retransmit. The back-off algorithm is designed to increase the average wait period before retransmission if more and more consecutive collisions are encountered. When the channel experiences heavy traffic from multiple CDPD subscribers, this will result in more collisions, which can cause longer delays in accessing the channel and successfully transmitting data.

DSMA–CD is very similar to Ethernet’s carrier sense multiple access with collision detection (CSMA–CD) approach for transmitting data on a shared channel.

Applications for CDPD include the following:

- Transportation: Dispatch and vehicle routing information to drivers, fleet managers, and customers.
- Mobile Professional: E-mail, Web, corporate application, and database access from anywhere
- Public Safety: Real-time access to critical information for law enforcement officers and fire fighters
- Health Care: Home patient monitoring, remote medical support, and on-site prescription orders
- Field Sales: Customer and sales order database access on the road
- Field Service: Work orders, service order details, customer and equipment histories, and product information to service employees in the field
- Telemetry: Monitor and control remote equipment (e.g., utility meters, vending machines, alarms, office equipment)
- Financial News: Up-to-date financial data, including stock quotes and analysis, bank rates, industry announcements, etc.
- Financial Transactions: Credit-card verification and other point-of-sale applications

CDPD offers many advantages over circuit-switched data, which includeIP–based and open specification (no royalties or licenses), full duplex operation, raw transmission speeds of 19.2 kbps, and an “always-on” connection, since once a CDPD modem is registered with the network, the user is connected indefinitely. In a properly designed and operating network handoffs do not affect this connection. All of this can be accomplished in a fairly cost-effective manner, as the additional infrastructure is minimal in comparison to other data services.

The primary disadvantages of CDPD are that the actual transmission speed is limited to 9.6 kbps due to protocol formatting, and real network speeds are closer to 4 or 5 kbps. In addition, to utilize CDPD in its most effective manner, dedicated channels should be used, eliminating capacity on the voice network. CDPD is also subject to co-channel and adjacent channel interference and remains highly dependant on existing cellular design for proper operation.

General Packet Radio Service

The next step in the evolution of data in the wireless world is the utilization of packet-switched services in what is known as 2.5G offerings. First to be examined is GPRS. GPRS supplements today’s circuit-switched data services and SMSs.

Like CDPD, GPRS involves overlaying a packet-based air interface on the existing circuit-switched network, in this case a GSM network. The GPRS system is implemented in such a way as to allow network operators to add a minimal number of new infrastructure nodes and upgrade software on some existing network elements. GPRS, as other packet networks, splits information into separate "packets," then transmits and reassembles them at the receiving end. The Internet is the most well known example of a packet data network.

Packet switching offers efficient use of the available spectrum, as GPRS radio resources are used only when users are actually sending or receiving data. With GPRS, rather than dedicating a radio channel to a mobile data user for a fixed period of time, the available radio resource can be shared between several users. As a result, a larger numbers of GPRS users can potentially share the same bandwidth and be served from a single cell. The actual number of users that a GPRS network can serve is dependant on the length or size of the packets sent and the constraints of the carrier's infrastructure as it relates to total spectrum made available to the customer.

GPRS facilitates mobile Internet functionality by allowing interworking between the existing Internet and the GPRS network. Any service used over the fixed Internet today, including file transfer protocol (FTP), Web browsing, e-mail, and telnet, will also be available over the mobile network via GPRS. As it uses the same protocols, the GPRS network will actually function as a sub-network of the Internet. It should also be noted that GPRS is not only designed to be deployed on mobile networks that are based on the GSM standard—the IS-136 TDMA standard will also support GPRS.

GPRS theoretical maximum speeds have been rated up to 171.2 kbps, with GPRS using all eight timeslots at the same time. This is significantly faster than the data transmission speeds possible over today's fixed telecommunications networks and current circuit-switched data services.

Also, as with CDPD, GPRS offers the customer an "always-on" connection and facilitates several new applications not previously available over GSM networks. Other new applications for GPRS, profiled later, include file transfer and home automation—the ability to remotely access and control in-house appliances and machines. GPRS allows network operators to maximize the use of their network resources in a dynamic and flexible manner along with user access to resources and revenues.

Enabling GPRS on a GSM network requires the addition of two core modules, the gateway GPRS service node (GGSN) and the serving GPRS service node (SGSN). The GGSN acts as a gateway between the GPRS network and public data networks such as IP. GGSNs also connect to other GPRS networks to facilitate GPRS roaming. The SGSN provides packet routing to and from the SGSN service area for all users in that area. In addition to these nodes, some other technical changes need to be added to a GSM network to implement a GPRS service. These include the addition of packet control units, mobility management to locate the GPRS mobile station, a new air interface for packet traffic, new security features such as ciphering, and new GPRS–specific signaling.

While we can see that GPRS is a valuable new mobile data option, offering major improvements in spectral efficiency, capability, and functionality, it is important to note that there are some limitations with GPRS. GPRS will have an impact on a network's existing capacity. As with any cellular network, there are limited radio resources that can be deployed. The extent of the impact to the network depends upon the number of timeslots, if any, that are reserved for exclusive use of GPRS. One positive aspect of GPRS is that it dynamically manages channel allocation, which facilitates a reduction in peak time signaling channel loading. In addition, achieving the maximum GPRS data transmission speed of 172.2 kbps would require a single user occupying all eight timeslots without any error protection. It is highly improbable that a network operator will allow all timeslots to be used by a single GPRS user, therefore actual customer speeds will be significantly less than the theoretical maximum. Typical single user throughput is likely to be 56 kbps in first-phase GPRS. With the phased implementation of GPRS, higher bit rates will become available on an incremental basis, reaching realistic speeds of up to 112 kbps as GPRS Phase 2/EDGE begins to emerge.

Enhanced Data Rates for GSM Evolution

The beginning of 3G evolution includes EDGE. EDGE is a 3G technology that can be introduced into both GSM and TDMA networks. EDGE will support speeds up to 384 kbps in wide areas and at considerably higher speeds within existing radio spectrum bands. EDGE provides significantly higher capacity and the same "always-on" connectivity that customers enjoy with CDPD and GPRS.

For carriers, EDGE offers a relatively cost-effective means for providing 3G services within existing spectrum. EDGE will provide the ability for carriers to offer single terminal roaming between TDMA and GSM networks in all frequency bands worldwide. EDGE is widely regarded as a key factor in access and service convergence between the GSM and TDMA worlds, leading to seamless global roaming with advanced mobile Internet services.

The first step in the deployment of EDGE will be the overlay of an enhanced core network based on GPRS technology to handle data transport within the network. The GPRS core network is a common standard for use in GSM, TDMA, and UMTS networks. Adding GPRS, an intermediate step in the evolution to 3G services, involves new and updated switching and network nodes in the core network. Further migration to EDGE usually involves a software upgrade to the radio base-station equipment. International standards for EDGE and GPRS require roaming between the two types of networks, therefore users with GPRS terminals will be able to use them on EDGE networks and vice versa.

The benefits of EDGE to the consumer are the higher bandwidths, enabling high-speed access to the Internet, multimedia and data services, and the "always-on" connections. Advantages to the carrier include a cost-effective way to migrate an existing GSM or TDMA wireless network to 3G without requiring the acquisition of new radio spectrum. EDGE also increases the overall network capacity as a result of its spectrum-efficient nature utilizing packet data transmission. EDGE offers as much as three times the capacity of GSM with GPRS and as much as seven times the capacity of TDMA networks.

Universal Mobile Telephone Service/WCDMA

The 3G migration path will also include UMTS, commonly referred to as WCDMA. WCDMA is a 3G technology designed to support high-speed multimedia services such as video, Internet, and videoconferencing.

The UMTS standard will interoperate with EDGE 3G deployments for both TDMA and GSM operators and provide transparent high-speed wireless data services to customers around the world. EDGE and UMTS share the same core network and are supported by the third-generation partnership project (3GPP)

A high-level overview of the upcoming evolution in data rates enabled via 3G WCDMA technology is subsequently quoted and can be found at www.mobilestreams.com.

3G Data Rates
The International Telecommunications Union (ITU) has laid down some indicative minimum requirements for the data speeds that the IMT2000 standards must support. These requirements are defined according to the degree of mobility involved when the 3G call is being made. As such, the data rate that will available over 3G will depend upon the environment in which the call is being made:

- *High Mobility*
  144 kbps for rural outdoor mobile use. This data rate is available for environments in which the 3G user is traveling more than 120 kilometers per hour in outdoor environments. Let us hope that the 3G user is in a train and not driving along and trying to use their 3G terminal at such speeds.
- *Full Mobility*
  384 kbps for pedestrian users traveling less than 120 kilometers per hour in urban outdoor environments
- *Limited Mobility*
  At least 2 Mbps with low mobility (less than 10 kilometers per hour) in stationary indoor and short-range outdoor environments. These kinds of maximum data rates are often talked about when illustrating that the potential for 3G technology will therefore only be available in stationary indoor environments.

In addition to the network equipment and protocols that will be needed to enable WCDMA, new cell planning methods will be required to support the 3G network rollouts. Basically, more 3G base stations will be required to provide coverage comparable to that which exists on current 2G networks.

The following information courtesy of http://www.ericsson.com/wcdma/img/WhitePaper_wcdma_ran.pdf provides additional insight into the workings of WCDMA:

> Code Division Multiple Access (CDMA) is a multiple access technology where the users are separated by unique codes, which means that all users can use the same frequency and transmit at the same time. With the fast development in signal processing, it has become feasible to use the technology for wireless communication, also referred to as WCDMA and CDMA2000. In CDMA, a narrow band signal is multiplied by a spreading signal (which is a pseudo-noise code sequence) with a higher rate than the data rate of the message. The resultant signal appears as seemingly random, but if the intended recipient has the right code, this process is reversed and the original narrow band signal is extracted. Use of unique codes means that the same frequency is repeated in all cells, which is commonly referred to as a frequency re-use of 1. WCDMA is a step further in the CDMA technology. It uses a chip rate of 3.84 Mcps, which is about three times higher than the chip rate of CDMA2000 (1.22 Mcps).
>
> The main benefits of a wideband carrier with a higher chiprate are:
>
> - Support for higher bit rates
> - Higher spectrum efficiency thanks to improved trunking efficiency (i.e. a better statistical averaging)
> - Better coverage as the frequency diversity is improved
>
> Further, experience from second-generation systems like GSM and cdmaOne has enabled improvements to be incorporated in WCDMA. Focus has also been put on ensuring that as much as possible of WCDMA operators' investments in GSM equipment can be re-used. Examples are the re-use and evolution of the core network, the focus on co-siting and the support of GSM handover. In order to use GSM handover the subscribers need dual mode handsets.

3G encompasses many aspects of service to the wireless customer and includes volumes of information and research. A high-level view of 3G includes IP–based data and voice communications over a packet data service node (PDSN) as opposed to the existing PSTN, which is ill equipped to provide the speed and accuracy needed for future wireless applications. Current wireless systems utilize the PSTN for both voice and data services. A typical call path today would include base transceiver stations, located at the tower sites; base-station controllers, located adjacent to mobile switching centers; and common fiber or copper transmission paths to the PSTN. Wireless data integrates the interworking function (IWF) into the path to create circuit-switched data calls for fax, paging, or data transfer and now include network interface and server devices to create a path to the Internet or VPNs.

3G technology will enable the current air interface to increase available bandwidth from 14.4 kbps to 144.4 kbps in its first release, scheduled for a phased implementation beginning in 2002. Speeds of 2 Mbps will be achieved in future applications of 3G. In addition to eliminating the current air-interface restrictions, 3G will incorporate packet switching, as opposed to circuit switching, for more capability in the data arena. PDSNs will be added to the call path for 3G data traffic and will utilize IP–based systems for routing and transfer of voice, data, and video services from a common device. An example of a 3G personal communications service (PCS) network can be found in *Figure 3*.

FIGURE 3

3G Network Configuration

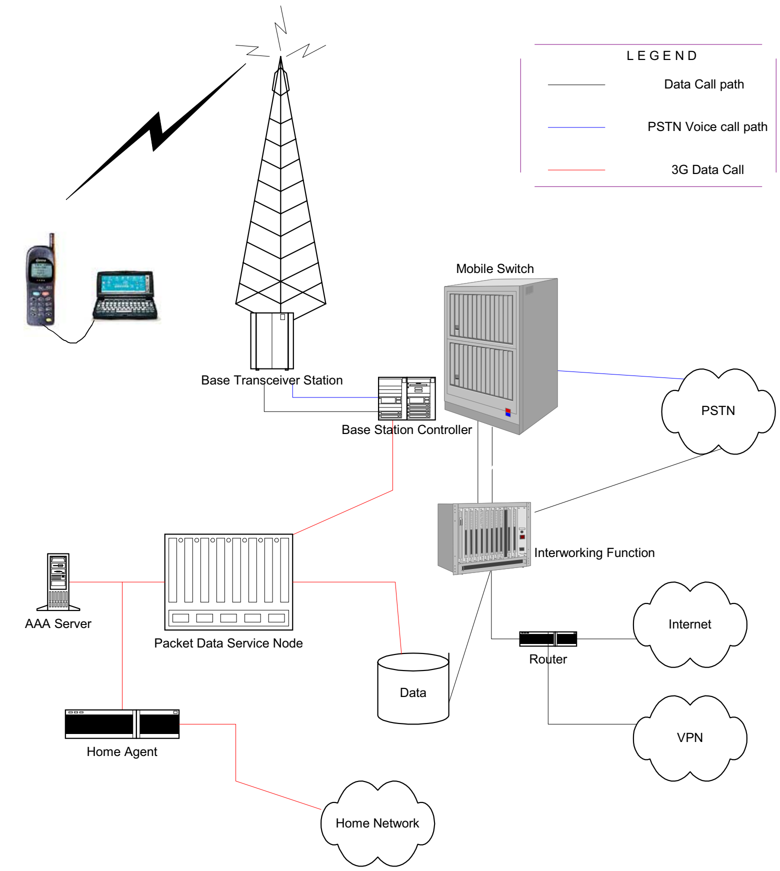

In the very near future, it will be possible for the customer to utilize a single handheld device to view Moving Pictures Experts Group (MPEG) files, make voice calls, receive e-mail and faxes, and surf the Internet simultaneously. In addition to providing expanded services, this IP–based system will allow for better network monitoring and troubleshooting, as well as provide billing and services in a very surgical manner. Another advantage offered by 3G is that bandwidth will be allocated on an "as needed" basis. According to Maria Palamara and Lucia Sellers of Lucent Technologies, "When a user makes a voice call, the system automatically allocates eight kilobytes per second. When the user wants to video conference or surf the Internet, the system automatically allocates more bandwidth" (http://www.lucent.com/wireless/products/solutions/tgwire.html).

3G will incorporate current CDMA, TDMA, and GSM platforms into upgraded, standards-based, worldwide deployable systems that provide the bandwidth needed by both business and home users. Palamara and Sellers sum up the potential of this technology as follows, "The third generation takes wireless communications to new levels, making it even more competitive with wire line communications in terms of cost and services. Just consider how remarkably the wire line communications marketplace has changed in the past few years. With the explosion of the Internet, phone lines now carry more minutes of data traffic than voice traffic. It is only a matter of time and technology before this trend hits the wireless industry and the third generation technology will make it possible to meet the high-speed data demands of any consumer" (http://www.lucent.com/wireless/products/solutions/tgwire.html).

3G will require a number of hardware and software upgrades, and in most cases will involve "forklift" upgrades to existing equipment. In addition, as with most new wireless technologies, the effective footprint of each base station will become smaller, as the reliance on clear, error-free data streams becomes essential. With a voice call, a missed bit or two equates to static or a missed letter or word, easily compensated for and still understandable to the callers. With a data call, that same missing bit or two results in a break in the data stream that can be interpreted by the receiving device as either an end to the message, resulting in a dropped call, or a change in the context of the entire data stream, which could generate spurious information. Either results in a complaint call from the customer. The solution is to provide better, contiguous coverage through more sites and shorter distances between them. This growth potential alone will keep infrastructure providers busy through 2005.

According to Peter Benson, a senior consultant with EMS-Global:

> The truth is that wireless data transfer isn't particularly fast or reliable. It currently runs at just 9.6 kilobits to 14.4 kilobits per second—a whole lot slower than the 56 kilobits per second speed of that modem hooked up to your PC at home that you complain is already too pokey. That's why the wireless companies are virtually giving the service away. Sprint PCS (Personal Communication Service, for example, will let you sign up for its wireless Internet service as a $10 a month add-on to its regular PCS service.
>
> But you know what? The wireless industry also isn't spending anything significant to provide this toy. Nobody's upgrading their network to deliver stock quotes to a few thousand customers. If somebody buys a more expensive Internet phone instead of a basic model from Motorola (MOT) or Nokia (NOK), that's gravy. The $10 a month is a little incremental revenue to pay for some of the marketing campaigns.
>
> ***Getting ready for the big rollout***
> All that the wireless companies are doing right now is getting the market ready for the rollout of the next generation of wireless broadband technology—or maybe it's more accurate to say technologies. In the time-honored tradition of the wireless industry, these companies are backing two competing standards:
>
> CDMA (Code Division Multiple Access) and General Packet Radio Service.
>
> Both CDMA and General Packet Radio Service offer really big increases in speed over the current 9.6 kilobits-per-second rate. According to Michael Murphy, editor of the California Technology Stock Letter, General Packet Radio Service promises 100 kilobits per second—about 10 times faster than the current wireless standard. The new CDMA standard will be even faster at 128 kilobits per second. That's about twice as fast as the current 56-kilobit home PC modem.
>
> That speed probably isn't fast enough, however. You know how long it can take to load an Internet page at home. No one will stand around in the street waiting that long. So these new technologies play tricks to make the Internet faster. For example, Nokia and other companies backing the General Packet Radio Service standard are also pushing a technology called Wireless Access Protocol, or WAP. A WAP server—and Nokia is currently giving these away—automatically rewrites a Web page from the HTML language into the simpler WAP language to speed up delivery.
>
> Wireless data speeds will continue to increase with time, but the development path for this technology is pretty clear. The next generation will be faster than the standard PC modem operating over standard copper wires, and about as fast as ISDN. (Remember that technology?) It will still lag the 640 kilobit speed of DSL, but even at 100 to 128 kilobits, wireless is in the game (www.ems-global.com).

Conclusion

As we can now see, based on the information and technological growth both documented and predicted, the entrepreneurial style that has defined wireless from the onset answers the question asked in the introduction of this paper. The long answer comprised the content of this paper and included an examination of wireless data technology from current HSCSD to early packet data and into advanced, high-speed packet-switched networks. The short answer is that the growth seen in the last decade of the 20th century is most definitely NOT "all there is." In fact, all indications are that the best is yet come and that the world of wireless data will continue to be an exciting, evolving place in which to live and work.

References

1. http://www.ericsson.com/wcdma/img/ WhitePaper_wcdma_ran.pdf
2. http://www.lucent.com/wireless/products/solutions/tgwire.html
3. www.allnetdevices.com
4. www.cdg.org
5. www.ems-global.com
6. www.gsmdata.com
7. www.ieee.org
8. www.mobile3G.com
9. www.mobilegprs.com
10. www.mobilestreams.com
11. www.msn.com
12. www.uwcc.org

Understanding Multimedia Messaging

Gideon Fostick
Director, Multimedia Messaging
Comverse

Chen Porat
Director, Multimedia Messaging
Comverse

Profitable operation of next-generation cellular networks poses a significant challenge for mobile network operators. Because the deployment of these networks requires substantial financial outlays, one of the greatest concerns for mobile operators is how to recoup their massive investment in network infrastructure. Operators need to offer end users enhanced communications services that will maximize return on investment (ROI) and strengthen market-share in this increasingly competitive marketplace. Designed for both 2.5G and 3G networks, and capitalizing on the Internet protocol (IP) connectivity and increased bandwidth of these advanced networks, multimedia messaging service (MMS) has the potential to meet this need.

Overview

Multimedia messaging introduces considerable enhancements to current mobile person-to-person messaging services—enhancements that promise to enrich the end-user experience. Imagine sending or receiving your stock portfolio summary with daily charts, a birthday card for a faraway friend, a promotional clip for a potential client, and new baby footage for Grandma—all via your mobile phone or personal digital assistant (PDA). MMS augments traditional text- and voice-based messages with any combination of rich-media elements, such as photos, images, and streaming audio, and video. MMS builds upon the short messaging service (SMS) user experience model by adding new capabilities easily understood by consumers. Given the impressive popularity and proven profitability of SMS, MMS is expected to experience rapid end-user adoption and market penetration.

An international in-depth research study, commissioned by Comverse and completed in January 2001, revealed that an overwhelming 80 percent of mobile phone users would enthusiastically embrace MMSs as soon as they become available. The study surveyed three key mobile phone user markets—residential consumers, business consumers, and young (15- to 16-year-old) consumers—in London, Frankfurt, Madrid, Milan, New York, Paris, Seoul, Stockholm, and Tel Aviv.

The results of this study provide compelling evidence for service providers to exploit the opportunity of 2.5G networks and initiate an MMS revenue stream without having to wait for the advent of 3G. Their introduction to launching MMS in 2.5G environments will then enable them to migrate seamlessly to more sophisticated 3G services. Although the basic function of MMS is person-to-person messaging, it should also be viewed as a powerful enabling environment that allows the provision of additional services, such as information and entertainment services, leveraging the capability for mobile multimedia content delivery.

To take advantage of this market opportunity, network operators must fully understand the link between the end user's desired experience and the advanced technology behind MMS. End users, now accustomed to using their mobile phone for non-voice communications, will expect Internet-like multimedia messaging capabilities over mobile networks. If MMS is to leverage the success of SMS, the MMS user experience must replicate the SMS user experience—immediate, fun, and easy to use. Though the market perceives MMS as the natural evolution of SMS, MMS requires far more advanced underlying technologies. MMS breaks from SMS's native "old world" circuit-switched environment, operating instead as an IP–based solution for dealing with multimedia files of varying sizes over multiple devices. To deliver true MMS functionality while preserving the SMS user experience, operators must build MMS systems that meet three critical requirements: the deployment of scalable and open platforms, support of legacy handsets with limited MMS capabilities, and provision of content and media to stimulate MMS service usage.

The Technology

MMS versus SMS

The only similarity between MMS and SMS is in the user experience—the underlying technology is completely different. SMS is based on traditional telephony protocols, such as signaling system 7 (SS7). MMS, on the other hand, is firmly rooted in the realm of IP messaging—a concept as different from SMS technology as the mobile phone is from the traditional fixed line telephone. To mirror the SMS user experience, MMS must utilize a system of rules that take into account user preferences, handset capability, and available network resources.

SMS has gained widespread popularity in recent years, most notably in Europe and parts of Asia, but its technology does not extend to the MMS arena. SMS is strictly text-based and limits messages to 160 characters. SMS uses the signaling channel between the handset and the base station to send and receive messages. This channel was selected because it is the only packet data channel that always remains "open," but its low speed and low capacity create significant constraints on traffic. The transmission of graphic and video data demands massive bandwidth. An updated version of SMS, known as enhanced messaging services (EMSs), uses a string of concatenated SMS messages to send larger multimedia content files. Even EMS, however, would heavily tax limited signaling network resources.

The different system design requirements and conceptual paradigm shift in moving from SMS to MMS are illustrated in *Table 1*.

Because of the technology differences described, major vendors of SMS have had to build a completely new architecture to handle MMS—without reusing proprietary SMSC technology. In MMS, legacy SMSC is retained solely for the notification stage of MMS transfer.

MMS versus Mobile E-Mail

Although MMS has technology roots in the e-mail domain, it has significant and distinctive characteristics that differentiate it from mobile e-mail services and enable it to provide users with a superior mobile messaging experience. The main differentiating factors are as follows:

- ***Push:*** MMS supports push and streaming as delivery methods, using store-and-forward architecture. Push is an important factor for the end-user experience: Users will receive an intriguing notification of new multimedia messaging, followed by an immediate view of the content, because the push mechanism automatically downloads the content before notifying the user. This quick and satisfying experience was a key factor in SMS success.

- ***Client-Server Negotiation Capability:*** The MMSC detects the multimedia capabilities of the recipient's device and performs comprehensive adaptation measures on the message content to ensure optimal display.

- ***Enhanced Multimedia Presentation:*** MMS minimizes the need to use attachments. Text and images can be presented in a synchronized manner as part of the message body while other audio and video elements are incorporated in the message body by using universal resource locator (URL) references. On MMS clients that support synchronized multimedia integration language (SMIL), all media elements can be presented in a synchronized manner.

Building an MMS Solution

In selecting the right technology and architecture to develop a robust MMS system, operators must consider the experience and legacy of the various vendors as well as the inherent technology, architecture, and management support. A system should incorporate the following:

- ***Strict Adherence to MMS Standards:*** The MMSC must be fully compliant with the third-generation partnership program (3GPP) MMS model as specified in the standards specifications TS22.140 (service definition) and TS23.140 (functional description), and with the WAP Forum MMS encapsulation standards (WAP 206 through 209). Adherence to these standards is crucial to ensure interoperability between the MMSCs and MMS handsets provided by all vendors.

- ***IP Store-and-Forward Technology:*** The proven IP store-and-forward structure embodies technology for high-volume real-time multimedia processing for effective content transcoding and streaming.

- ***Flexibility:*** MMSCs must accommodate different message processing flows, defined by operator policy, subscriber preference, client operating system, or message characteristics such as type or size.

- ***Openness:*** In the MMS world, content and technical partnerships will be crucial for creating business models and service platforms. MMSCs must support these models by offering highly open architectures and dedicated interfaces based on open IP protocols. This openness will attract third-party value-added service providers.

- ***Protection of Current Infrastructure Investments:*** A substantial number of operating companies throughout the world are already deeply invested in IP–messaging infrastructure, including equipment for connectivity, storage, security, load balancing, and management. This infrastructure supports existing systems, such as unified messaging. Reuse of existing system components and infrastructure will optimize investments and save footprint, which in turn will translate into a highly competitive user service.

- ***Performance:*** The processing of multimedia traffic will result in dramatic increases in both the volume and the size of the messages. MMS infrastructure must be designed from the outset to deal with these enormous throughput requirements, expected to reach more than 1,000 multimedia messages per second.

- ***System Management:*** Managing a large and functionally rich distributed MMS system requires extensive facilities. Sophisticated management and billing tools—multi-site, Web, or graphical user interface (GUI)–based—must be made available as a part of the MMS solution and must be configurable for distributed MMS networks. Advanced tools for system monitoring and statistics analysis are important for optimizing the MMS service.

- ***Billing:*** MMS must conform to the maxim, "Bill it or kill it." Extremely flexible billing solutions, for both prepaid and postpaid subscribers, are necessary to accommodate evolving billing models, as operators learn how to best charge for MMS services. Billing solutions should handle straightforward models, such as event-driven billing, based on volume, content, or message type, as well as complex billing situations, such as reverse- and reply-charging and revenue collection and sharing for third-party service providers.

- ***Service Flow Flexibility:*** MMS systems will need to adapt dynamic subscriber behavior to service models. For this purpose, leading MMS systems are designed with flexible service behavior definition tools, such as those based on the extensible markup language (XML)

TABLE 1

MMS Functionality

Functionality	Short Message Service Center (SMSC)	Multimedia Message Services Center (MMSC)
Protocol	• Message to handset via SS7, MAP	• IP–based protocols, WAP Encapsulation, SMTP
Message characteristics	• Text messages/icons • Small: 160 bytes • Fixed size • Simple format	• Multimedia messages • Large: Megabyte levels • Varying sizes • Complex MIME format, nested and multipart messages
	This profound disparity in message characteristics dictates that the two platforms require totally different data structures and handling processes.	
Message delivery	• One-stage delivery: message push • Single SMSC handles end-to-end delivery • No adaptation of content	• Two-stage delivery: Notification push followed by content retrieval after an indefinite time period • Two MMSCs (at sender and recipient networks) cooperate for delivery. Raises discovery, routing, and interoperability issues • Content adaptation and transcoding—can require heavy processing
	These different message processes break the MMS delivery mode away from SMS and require processing overhead, which SMS platforms do not support.	
Facilities for external applications	• Requires multiple and specific protocols/solutions: SMPP/CIMD2/UCP	• Easily implemented using standard Internet protocols, e.g., SMTP, XML/HTTP
Capacity planning	• Load balancing is not easily applicable to SS7, thus planning for heavy loads is based on highly specialized proprietary platforms • Closed standards implemented by a few specialized vendors	• Use of Internet SMTP technologies ensures scalability, including load sharing and balancing • SMTP technologies are widely sourced by many vendors and proven scalable by prevalent use in Internet systems
Streaming	• Not possible	• Possible and desirable

service description scheme, which allows the addition of service components and applications as an integral part of the process. In contrast, SMS systems are closed and proprietary by nature and have no need or facilities for such dynamic service forming.

- ***Security:*** MMS is implemented in the open IP world, creating the potential for the sort of malicious acts, spamming, and unsolicited material that plague e-mail and the Internet. Both operators' and users' concerns make this a high-priority issue.

- ***Scalability:*** MMS systems should be designed with an eye to the future, enabling many options rather than being locked in to a single inflexible program. As the market matures, savvy customers will request MMS enhancements, such as increased storage and handling of messages, multi-access, and unified mailbox integration.

Expanding MMS Horizons

Multimedia Album
As subscribers adopt MMS, they will require specialized resources to help them compose fun and expressive multimedia messages. One such resource should be the Multimedia Album. This is an on-line repository of "canned" multimedia elements such as pictures and video clips—similar to the concept of clip-art files—from which a user can select professionally created images to enhance their messages. The Multimedia Album should support a variety of multimedia content types and should be accessible by Web, WAP, and standard MMS handsets. The Multimedia Album should also accept content provided by the subscribers themselves, by the operator, or by third parties. Automatic and periodic content updates should be supported while maintaining the security of the album content and interfaces.

Safe Storage
Subscribers are sure to demand safe and unlimited storage for multimedia content. However, handset memory will always be limited and may lag far behind demand as people start using handsets to view rich media such as pictures and video. As MMS usage becomes common, content-storage security will become a very important issue. As distinct from SMS, multimedia messages are bigger and more precious: "I don't want to lose valuable pictures (of family, friends, or special occasions) if I lose my handset or drop it on the floor." This suggests a new and valuable service for operators to provide network-based storage. By providing this service, operators will become trusted agents for storing users' valuable content, thereby enhancing relationships with subscribers and reducing churn.

Overcoming the Handset Challenge

A key impediment to mass-market acceptance of MMS is the expected delay in the production of MMS handsets that incorporate a dedicated client to support the standard. Such handsets may gradually trickle into the market, too slow to meet the demand of consumers who are intrigued and eager to engage in the MMS experience. However, by the innovative use of existing technology, a taste of MMS will be available to virally stimulate the market far in advance of the mass-market penetration of the devices designed to deliver it.

The lesson drawn from SMS success is that, as with all communication technologies, demand and value of MMS will only increase as the technology becomes more widely used. As an example, if one person has a telephone, he or she has nothing, but one million people with telephones have valuable communication tools. In Europe and Asia, SMS took off because most people had SMS–capable handsets. In the United States, where only a portion of the population had SMS–capable handsets, SMS did not experience the same success.

Operators can overcome the handset challenge by offering MMS systems that support legacy devices. An MMS system must be intelligent enough to determine the capabilities of the target device and adapt the content delivery as necessary. The device can be anything from a legacy 2G phone (e.g., Nokia 6110), a wireless access protocol (WAP) phone using circuit-switched service (e.g., Motorola Timeport), a GPRS WAP device (e.g., Siemens S45), a PDA (e.g., Cassiopeia) or a fully compliant MMS device (e.g., Ericsson T68).

Subscribers with legacy handsets will then be able to receive and enjoy multimedia messages and, to some extent, send them as well. Legacy devices will be able to access MMS messages via adapters to widely accepted clients—mobile and personal computer (PC) Web browsers, WAP browsers, and e-mail clients (IMAP4/POP3). This MMS experience may be more limited than the full-blown service, and sending messages may not be a streamlined process, but the initial experience will be available. It is important to note that full billing capability should be preserved; operators must be able to charge for MMS use by legacy devices, both for prepaid and postpaid subscribers.

The important implication is that those people who have purchased dedicated MMS handsets can immediately start sending multimedia messages to friends and colleagues with legacy devices. Introducing MMS to the market in this way creates the potential for a dynamic viral effect, whereby early adopters start using MMS with each other and toward friends with legacy devices, thereby stimulating and encouraging their friends to join the MMS "party" and purchase their own MMS–capable handsets.

By virally introducing the MMS experience and giving users with legacy handsets their first taste of MMS, it is estimated that uptake of MMS could be accelerated by 18 to 24 months.

Summary

MMS is poised to initiate a totally new form of communication. It may well be as big as, if not bigger than, SMS, and while it is tempting to follow the SMS model, the axioms are profoundly different and require a conceptual quantum leap.

The message characteristics and delivery flow are remarkably different, thus dictating that the MMSC messaging platform requires totally different underlying technology from that of the SMSC. Despite the anticipated delay in MMS–capable handsets reaching full market penetration, MMS can take off right now, allowing early adopters with MMS handsets to send messages to friends with legacy devices, thereby providing a viral dynamic for introducing people to the taste and fun of MMS. This will accelerate the uptake of MMS by up to 24 months, providing early entrant operators with a significant market advantage.

MMS is available here and now, in 2.5G, and operators need to get ready.

Prospects of Secure Real-Time Video Transmission over CDPD Networks

Forouzan Golshani

Professor, Computer Science Department
Arizona State University

Gamze Seckin

Research Engineer, Media Networking Group
Luxxon Corporation

Farid Hatefi

Solution Architect
Integrated Information Systems
Ph.D. Candidate
Arizona State University

Abstract

With the current channel limitations of cellular digital packet data (CDPD) networks and also the lack of bilateral authentication, it is difficult to meet the bandwidth needs of secure real-time video. Various error sources make it difficult to meet the quality of service (QoS) requirement within an acceptable bit-error rate and delay. Access point limitations and additional security measures add more traffic overhead. Adaptive control over the network traffic provides extra support for end-to-end QoS. Thus adaptive multimedia applications are at the core of wireless media streaming [12]. In this paper, we first show that CDPD is suitable for real-time video streaming by use of the adaptive mediator system [11]. Then we provide an end-to-end encryption and bilateral authentication technique based on X.509 certificate between the mobile units and base stations. The overhead traffic resulting from the additional encryption and authentication layer is one of the major challenges for secure real-time video transmission over CDPD.

1. Introduction

Wireless and wired networks have similar security requirements, namely authentication of who the other party is, securing the data as it travels from the handheld device to the destination host, and ensuring that the traffic hasn't been altered en route. However, wireless networks present some unique difficulties, including limitations on bandwidth, higher levels of latency, and instability of connections. Several techniques have been proposed to address these issues, more notably, wireless transport-layer security (WTLS), which functions similar to secure sockets layer (SSL), and connection-oriented security communications using traditional protocols such as Internet protocol security (IPSec) and secure shell (SSH).

Following the complete standardization of intelligent transportation systems (ITS) in 1993, full implementation of CDPD is possible. We use the work by Pereira [1] as the basis for our work for video transmission over CDPD. In [1], the performance of a minimal CDPD network is analyzed for different applications. The CDPD network is then deployed in a major metropolitan area (in Pennsylvania in 2002), where the infrastructure is shared by all wireless applications (meaning that no CDPD channel will be reserved exclusively for ITS).

A CDPD forward link channel operates in broadcast mode. Thus there is no contention involved, other than occasional retransmission with the possibility of some queuing. The performance challenge is on the reverse link, where many users compete for a single channel. The CDPD random access protocol conforms to a non-persistent digital sense multiple access with collision detection (DSMA–CD). The forward channel generally operates far below capacity. This allows for reducing the delay in the forward link to queuing delay. For the reverse channel, due to the inherent contention mechanism, only simulation can provide the necessary delay/throughput information.

The observed performance of CDPD networks confirms the adequacy of CDPD for wide-area wireless data service, accommodating even the most demanding and time-critical applications. Sending real-time video data over a CDPD network may require strong encryption of video traffic and authentication of all involved parties in the communication. The current CDPD specifications, R.1.1, have not provisioned a mechanism for bilateral authentication between the mobile end system (M–ES) and serving mobile data intermediate system (MD–IS), i.e., base stations. The available security requirements within CDPD networks address only one-way M–ES authentication by home MD–IS (also known as mobile home function, or MHF). In addition, there is no encryption capability provided between mobile serving function, or MSF (when the user is roaming outside of the home area), and MHF. The current encryption is not enough to prevent active attacks against CDPD traffic. In this paper, a new protocol, which provides bilateral authentication and packet encryption privacy, is proposed. Due to scalability and simplicity for generalization, this proposal may be seen as addressing the security aspects of an infrastructure for wireless environments.

2. Video over CDPD

Although 2.5G and 3G wireless solutions raise high expectations, their full deployment requires some time, and in the meanwhile work on 2G systems continues. According to estimations by GTE research [1], more than 60 percent of the market will belong to public wireless usage and non–ITS applications in 2002.

The GTE research also indicates that non–ITS data load will be 2.5 times the ITS load in the reverse direction and four times higher than that in the forward direction formulas 1 and 2. The reason for different treatment is that reverse direction traffic is mainly due to e-mail alone, while forward traffic is expected to be e-mail plus World Wide Web (WWW) access.

Reverse_Direction CDPD Channel → Non_ITS_Load=2.5 x ITS_Load (1)

Forward_Direction CDPD Channel → Non_ITS_Load=4 x ITS_Load (2)

The bit rate for a typical CIF (352x288) formatted uncompressed 10 fps video is around 12.2 megabits per second (Mbps), and around 36.5 Mbps for 30 fps video. Considering the limited bandwidth provided by CDPD networks, we have to limit the quality of video transmission. A simple multimedia application that requires the transmission of one-way video (SQCIF 128x96) with 10 fps and one-way telephone quality audio can enable a satisfactory visual implementation environment for CDPD networks. As known, an 8-kilohertz (kHz) sampling of telephony quality audio requires around 64 kilobits per second (kbps) of bandwidth. The overall uncompressed bit rate will be (983 kbps plus 64 kbps) around 1 Mbps. With a high-performance compression scheme, bit rate can be reduced to 10 kbps. Considering the maximum end-to-end delay requirement of 150 ms of meaningful video, the results reported in [1] indicate an average delay of 200 ms for more than 50 percent throughput for both ITS and non–ITS traffic, which prove CDPD to be suitable for video transmission.

The near future of video coding will be content based [2, 3, 4, 12], meaning each video stream will consist of several video objects as provisioned by Moving Pictures Experts group (MPEG)-4 [4]. Each video object will have different scalability levels, content information, bit-rate requirements, and QoS parameters. A source may have multiple video objects, each with multiple importance layers. The definition of video objects within a frame will be application-dependent. Each video object will have different characteristics and priorities. The advantage of using priorities and multiple objects is that the transmission system becomes more flexible for losses.

On the other hand, content adaptation and scalability is needed to allow wireless operators to support the increasingly diverse mobile wireless devices (personal digital assistants [PDAs], cell phones, communicators, laptops, etc.) with different video and audio capabilities (operating system [OS], bandwidth, display resolution, etc.) to respond to the fluctuating wireless bandwidth during a stream session, to support the constantly evolving wireless bandwidth (CDPD, IS-95B, CDMA2000 1x, general packet radio service [GPRS], wideband code division multiple access [WCDMA]) and to enable content to be created in a one format and played back in any of a number of formats and characteristics. Adaptive control over the Internet is providing extra support for end-to-end QoS. Thus adaptive multimedia applications are at the core of wireless media streaming [12].

3. Importance of Wireless Security

Some of the most publicly used wireless networks today suffer from architectural security flaws. Wireless (or wireless fidelity, or Wi-Fi) local-area networks (WLAN) based on 802.11b can be compromised with little effort. There are several security flaws in the wired equivalent privacy (WEP) algorithm, which is responsible for establishing a secure tunnel between the mobile user and the base station. Some of the major flaws in WEP are 1) passive attacks based on statistical analysis for decrypting the radio frequency (RF) traffic, 2) active attacks based on known plain text with the purpose of adding new packets to the traffic flow by an invalid mobile unit, or 3) dictionary attacks that help a perpetrator to conduct real-time decryption of traffic. The end result is an insecure network that provides a false sense of security for its users [14].

Another standard that has some security defects is wireless application protocol, or WAP [16]. Although WAP is using WTLS, which is a subset of SSL and transparent LAN service (TLS), it lacks provisions for providing full-scale security to its users. This is due to the changes that had to be made to TLS to make it more prominent for wireless mobile users. Because of these changes, WAP is now more susceptible to some attacks, such as chosen plain-text attack based on predicting the initialization vector, brute force attack against 40 bit key in data encryption standard (DES) algorithms, and the use of unencrypted and unauthenticated "warning messages" in between regular packets. These warning messages are issued by the base station, and they can be altered and resent to a mobile user by a perpetrator with devious purposes [15].

Although CDPD standard is one of the few wireless networks that is designed to operate and utilize several data communication security standards, it fails to meet some others, such as bilateral authentication, end-to-end data confidentiality, nonrepudiation, and traffic-flow confidentiality. The availability of inexpensive and sophisticated RF scanners that can be easily obtained (approximately the size of a credit card) by anyone, adds more emphasis to the importance of providing better security for a typical wireless network.

The most important determining factor for choosing a mobile network over others is how well it can create a real sense of security for its users. Also, the users must be assured that they are the only valid users of the channel.

3.1. Security Loopholes in CDPD Networks

CDPD specs are designed so that a central authentication process, located at the home service provider, known as the mobile home function (MHF), can authenticate all mobile end systems (M–ES). In addition, the wireless traffic on the air is encrypted for privacy of the messages between the M–ES and the MSF (when the user is roaming). However, as can be seen in *Figure 1*, the CDPD network does not provide any bilateral authentication between all parties. For instance, a mobile unit (M–ES) does not authenticate the service provider at the roaming area (MSF), nor vice versa. In other words, there is no bilateral authentication process between the involved parties in the aforementioned scenario. The provided security in CDPD networks blocks all of the static attacks on the radio traffic, but this does not protect the end user from dynamic attacks. A general security rule states that if the worth of information that can be sniffed off the air is much higher than the total cost of the attack, then such an attack is highly possible. It is possible to masquerade a serving MD–IS (service provider in the roaming area) in the CDPD network. As a result of that, a disguised serving MD–IS, after overpowering the real serving MD–IS RF signal strength, can obtain the end-user credentials. In such a case, when a mobile end user observes an acceptable signal strength from a serving MD–IS, it will participate in the handshake and key exchange procedure with the unauthentic serving MD–IS. After receiving the shared secret parameters from the mobile user, the fraudulent entity decrypts the messages and extracts the most important credentials from it, i.e., a unique network entity identification (NEI) number and a special sequence number known as shared historical record, or SHR. By now, not only does the disguised MD–IS have full access to the CDPD network, but the home service provider also will deny the real mobile user during the next authentication phase. This is due to the lack of knowledge of the latest SHR by the mobile end user, which is altered each time by the home MD–IS. The only entity that has the full knowledge about the latest SHR is the fraudulent MD–IS and mobile user who is using someone else's credentials [5,8].

There is another security problem that surfaces in CDPD networks when an attacker eavesdrops on the traffic between the serving and home MD–ISs. As can be seen in *Figure 1*, all of the traffic between the MSF and MHF is exchanged on the terrestrial line. There is no encryption or authentication mechanism between these two entities. The redirect request and redirect confirmed packets (RDR and RDC), which carry mobile end-user credentials, are exchanged in clear text between the MHF and MSF.

In the following section, a security mechanism is proposed that protects all the entities involved in the CDPD network. This mechanism provides a full end-to-end encryption as well as a strong authentication scheme. This is due to providing a bilateral authentication between M–ES and a serving MD–IS and home MD–IS [8].

3.2. Proposed Enhanced CDPD Security Algorithm

The new proposed approach takes advantage of a certificate authority (CA) system based on X.509v3 [9] standard. It provides end system authentication and a secure channel for data exchange through a combination of public key and private key algorithms. All messages between M–ES, MSF, and MHF are encrypted with RSA and RC4 [6,7] after verifying the authenticity of both the sender and the receiver.

Before reviewing the proposed security algorithm, it is essential to know that all involved entities in each session must already be registered with a CA and must own a certificate and also have access to the CA's public key. This provides a trust relationship between all entities. However, in the initial stage of communications between any two parties—e.g., a request for a connection and subsequent handshake process—there is no trust between the involved entities until each M–ES validates the authenticity of other M–ES's certificate.

FIGURE 1

Required Steps for M–ES Authentication and Message Encryption

STEP	M-ES		MSF		MHF
1		TEI REQ. →			
2		← TEI Assigned			
3		← $a, p, a^y \bmod p$ IKE			
4		→ $a^x \bmod p$ EKE			
5	$a^{xy} \bmod p$	Shared Secret Quantity	$a^{xy} \bmod p$		
6		ESH <NEI,SHR> → Encrypted Message	Decrypt	RDR <NEI,SHR> → Clear Text	M-ES Authentication
7		← ISC <NEI,SHR> Encrypted Message	Encrypt	← RDC <NEI,SHR>	

In the initial stage of the connection, the M–ES communicates with the MSF by sending a request protocol data unit (PDU) and its certificate for a new identification number or temporary entity identification (TEI). After authenticating the M–ES, the MSF in return sends the TEI, MSF certificate, and intermediate system challenge number (ICN), which is encrypted with MSF public key, to the M–ES. Upon receiving the reply, the M–ES uses the CA's public key to confirm the CA's digital signature. After confirmation, the M–ES extracts the MSF public key from the certificate and decrypts the ICN with its own private key. Then it generates two random prime numbers known as the master key (MK) and end system challenge number (ECN). Next, M–ES uses the MK to create two secret keys known as read and write keys, which are used in RC4 symmetric cryptography algorithm. In the next step, the M–ES sends a packet to the MSF consisting of the ICN, which is encrypted with the write key, ECN, and MK and encrypts the whole PDU with the MSF's public key. Upon receiving this "tuple," the MSF uses its own private key and decrypts the data and extracts the MK and ECN, and encrypted ICN. MSF uses the same algorithm that the M–ES used and creates two secret keys based on the MK (read and write secret keys). Next, the MSF uses the read key and decrypts the encrypted ICN, and confirms that this is the same ICN that it has created before. This insures that the MK is issued by the authentic M–ES. After confirming the ICN, the MSF uses its write key and encrypts the ECN, and sends it back to the M–ES. Upon receiving this message, the M–ES uses its read key and decrypts the message to confirm the ECN. This insures that the MSF is the one that issued the TEI. At this point the (A) interface (air channel) is secure between the M–ES and MSF, and the M–ES may send the end system hello (ESH) PDU. The same type of security handshake must be used between the MSF and MHF for securing the channel between them.

3.3. Analysis of the Algorithm

The enhanced security algorithm protects the visiting M–ES against a masquerading MSF. If a fraudulent MSF tries to masquerade its authenticity by sending a wrong certificate, the M–-ES will realize that as soon as is it reconfirms the CA's digital signature. It is extremely expensive (even impossible) to counterfeit the CA's digital signature, because the digital signature is the CA's encrypted hash value [6] of CA's encrypted hash value [10] of the contents of the certificate. CA uses its own private key to encrypt the hash value. Theoretically, the only entity that knows the CA's private key is the CA itself. Without the knowledge of the CA's private key, the fraudulent MSF must use a different pair of public/private key for encryption. If this happens, the confirmation of the CA's digital signature fails by the M–ES. If the fraudulent MSF captures the previous sessions' packets during the handshake and tries to replay them, the M–ES will accept the digital signature, but the masqueraded MSF wouldn't be able to extract the MK from the M–ES reply message. This causes the fraudulent MSF to not be able to create the read/write session's keys.

The enhanced security algorithm will also protect the MSF and MHF against static attacks during the exchange of RDR and RDC. Once the serving and the home MD–IS are both authenticated, they are the only entities that know the ses-

FIGURE 2

Required Steps for Enhanced Security in CDPD Networks

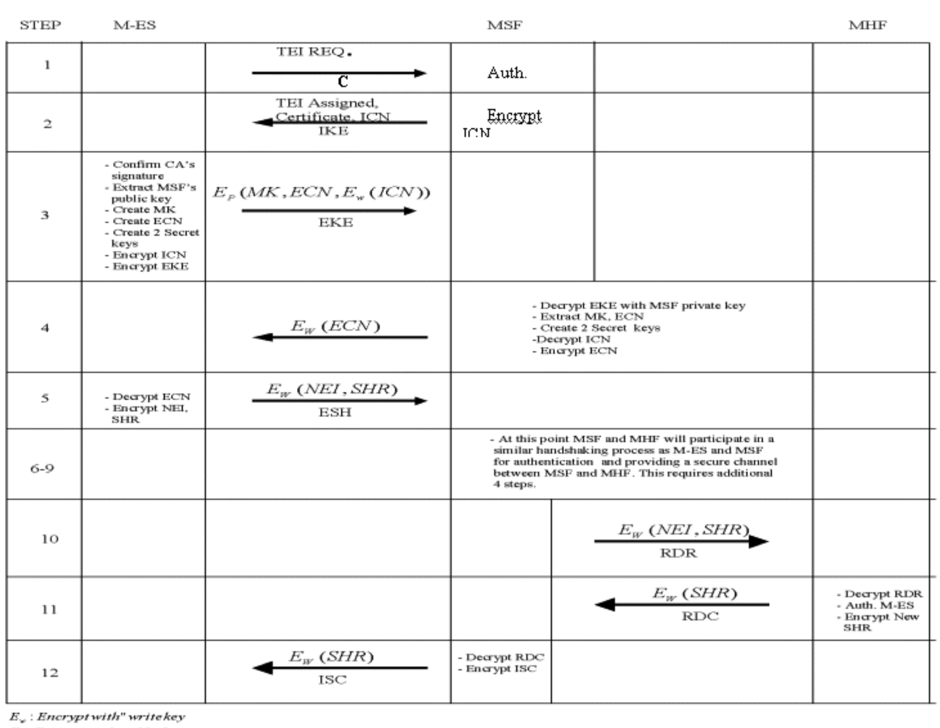

sion read/write key. This will create a fully secure channel between them. In case an eavesdropper tries to capture the ongoing packet between MSF and MHF, everything would be encrypted and meaningless.

4. Performance Considerations

The proposed additional security layer introduces some overhead, which impacts the performance of the system. In the following section, some measurements and comparisons are made in order to illustrate these impacts. These following measurements are based on the best condition in a CDPD network, which is behaving like a low-bandwidth wired network.

4.1. The Impact of Security on System Performance

Based on the research done in [17], a typical 60-bytes-long message needs additional 24 bytes (16 bytes for MD5, 4 bytes source time stamp, and 4 bytes for destination time stamp) to include data integrity and authentication information in the messages. This increases the original size of the message by 40 percent. Typically, this results in an increase of the service time for that message. If depicts the service rate of the CDPD system without the additional security layer, and is the service rate after adding the security layer, the following equation illustrates the service time increases due to a lower service rate after adding a security layer:

$$s = \frac{\mu}{\mu'}$$

Also, the real throughput of the system is the minimum value of for the different components (e.g., M–ES, MSF, MHF) in a CDPD network:

$$\mu_b = \min_{i \in [1,..n]}(\mu_i)$$

The statistics show [17] that the service rate for 100-bytes packets is about 1,900 packets per second. Packet authentication and authorization drops this service rate to 1,500 packets per second, implying a 20 percent increase in service time. On the other hand, the encryption alone drops the service rate to 630 packets per second (a 300 percent increase in service time).

Taking into account the packet retransmission time and transmission control protocol (TCP) congestion-control mechanism, the harsh wireless environment lowers the service time even more. To remedy some of these side effects, the TCP layer in a CDPD network must be optimized for such an environment.

5. Testing Environment

The initial set of tests for our research includes a feasibility study of video streaming over CDPD. To show that real-time video-streaming quality is acceptable over CDPD, we used Luxxon's adaptive mediator system [7] over a CDPD modem connection.

The mediator system is a real-time transport protocol (RTP)/real-time streaming protocol (RTSP)–based wireless streaming solution that takes in open standards–based MPEG video and audio, provides broadcast service through real-time video and audio transcoding, and can also perform off-line video and audio transcoding to create media-on-demand content for storage on file servers. The mediator system meets the challenges of wireless streaming by performing the adaptation, which is necessary to handle the diversity of devices, formats, and network characteristics. It provides bandwidth adaptability, end-device detection, and transcoding media to various bit rates and formats. Through these adaptability features, the mediator system supports various wireless bandwidths for 2G, 2.5G, and 3G networks, such as CDPD, IS-95B, CDMA2000 1x, general packet radio service (GPRS), and wideband code division multiple access (WCDMA).

Figure 3 shows a MPEG-4 video and GSMAMR audio clip with a frame rate of 5 fps and bit rate of 11 kbps. The CDPD connection was established through a UNIDEN Data-1000 modem with an average received signal strength indication (RSSI) of –76 dB with 600 mW antenna power output.

6. Conclusion

Any proposal related to an infrastructure for wireless and mobile networks must include an urgent discussion of how the so-called "gap in the WAP" can be plugged. This gap surfaces in a variety of ways in all other frameworks. Clearly, it is the device makers and the service providers that should provide the security protection. However, it is imperative for the consumer to understand the threats, be prepared with the necessary tools to combat them, and show readiness to deal with them should precautionary measures are shortcut.

FIGURE 3

Sample of Luxxon Player Streaming an MPEG-4 Clip at 11 kbps

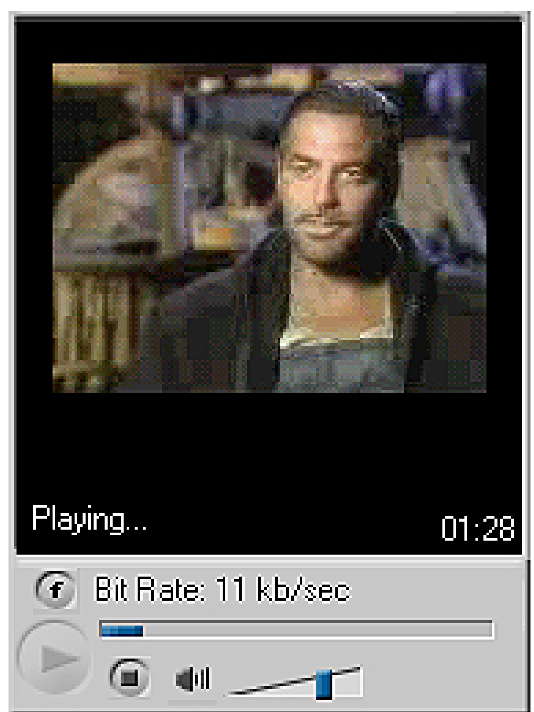

In this paper, we show that real-time video streaming is possible over CDPD networks. We suggest the usage of adaptive video over CDPD networks for better bandwidth utilization. To enable adaptive video streaming over CDPD, we have used the mediator system of Luxxon Corporation. X.509 v3 certificates and public key/private key encryption are used for further security enhancement of the video traffic. The X.509 certificates provide bilateral authentication for both M–ES and MHF. Also, the performance impact due to additional proposed security layer to the CDPD network also needed to be addressed. One possible solution to this problem would be an optimized indirect TCP (I–TCP) being implemented at the mobile database station (MDBS).

7. References

1. Pereira, J.M. and Fagen, D., "Personal, Indoor and Mobile Radio Communications," PIMRC'96, Vol. 2 (1996), 678–682.
2. Chan S., Wong E.W.M., and Ko K.T., "Fair Packet Discarding for Controlling ABR Traffic in ATM Networks", *IEEE Transactions on Communications*, Vol. 45, No. 8 (1997), 913–916.
3. Seckin G., "A Prioritized Congestion Control Method for ATM Networks," Proceedings of Third Symposium on Computer Networks – BAS'98 (Izmir, Turkey: June 1998), 80–87
4. Seckin G. and Golshani F., "A New Model for Multilayer Real-Time Video Transmission for ATM Networks," Proceedings of the Seventh Int'l. Conf. Comp. Commun. and Networks – ICCCN'98 (Lafayette, La.: October 12–15, 1998), 442–447.
5. Banan M., et al., *Internetwork Mobility: The CDPD Approach* (New Jersey: Prentice-Hall, 1996).
6. Stallings W., *Snmp, Snmpv2, Snmpv3, and Rmon 1 and 2*, (Addison Wesley, 1999).
7. Russell D. and Gangemi G.T., *Computer Security Basics* (California: O'Reilly & Associates, 1992).
8. Frankel Y., et al., "Security Issues in a CDPD Wireless Network," *IEEE Personal Communication* (1995).
9. ITU–T Standard, The Directory-Authentication Framework, CCITT Recommendation: X.509v3, 1988
10. Rivest R, The MD-5 Message-Digest Algorithm, Request for Comments 1321, Internet Activity Board (1992).
11. Luxxon Web site: www.luxxon.com
12. Seckin G. and Golshani F., "Real-Time Transmission of Multilayer Video over ATM Networks," *Computer Communications Journal*

Earning Revenue with Third-Generation Wireless: Killer Applications

Allen Groh
Director, International Trade and Regulatory Compliance
Ericsson

Sources of revenue in telecommunications and media are constantly changing—the result of ever-advancing technology. In the early days of long-distance service, charges for minutes of long distance were very high, due in large part to the cost and inefficiency of carrying analog signals long distances over copper wires. Advances in transmission technology from copper wires to microwave to satellites to fiber-optic cable and now to dense wavelength division multiplexing (DWDM), which dramatically increases the capacity of fibers, have made long-distance phone calls very inexpensive. Bandwidth has become a commodity, much the same way that long-distance service has become a commodity.

For third-generation (3G) wireless systems, financial analysts and service providers are asking, "Where will the revenue come from?" The cost of 3G wireless licenses in Europe has put the pressure on return on investment (ROI) for deploying 3G services. The answer is most likely in services offered via high-speed data networks that include voice traffic. The next question is "What will be the applications that will attract the most users and produce the most revenue?" The traditional source of wireless revenue—air time—is also becoming a commodity. Competitive pressure from as many as five wireless carriers in one metropolitan area helps to drive down the price of air time and thus decrease profitability for voice services. The introduction of data services will increase revenue only if a large number of customers subscribe. So far, wireless application protocol (WAP) offerings have had limited success. What is clear, though, is that significant add-on revenue can be earned by deploying data service along with voice only if enough customers will pay money for it

Killer Applications Defined

In the information age, killer applications have produced broad market appeal for consumer electronics in many forms. One of the first was the electronic calculator, which was one of the first consumer applications for newly invented integrated circuits. The electronic calculator, which was based on an early integrated circuit designed for missile guidance, replaced mechanical calculators and eventually the slide rule.

Advances in semiconductor design produced the first engineering calculator, the HP 35 followed by the HP 45 (both from Hewlett-Packard), which could do the job of the most sophisticated mechanical calculator faster and which cost 10 times less—about $400 versus $4,000 for an equivalent mechanical calculator. Soon, as the price came down and functionality increased, calculators made engineering slide rules obsolete (1975).

Perhaps the most famous killer app that launched the age of personal computers (PCs) was VisiCalc, an electronic spreadsheet program similar to today's Microsoft Excel. It literally saved a forest of trees in the form of paper spreadsheets and tons of rubber used in erasers. VisiCalc played another, much more important role. In the late 1970s, PCs were considered toys for hobbyists. Apple and Sinclair owned the market, but applications, where they existed at all, were more entertaining than useful. One could hardly make a case for PCs as a business tool, let alone for becoming an important economic driver of the 1980s and 1990s. VisiCalc, which was programmed to run on the Apple II computer, made the job of accountants infinitely simpler in preparing budgets and business plans. Just the idea that changing one number anywhere on a spreadsheet would result in automatic recalculation of the entire page revolutionized the financial planning segment of the business world. Sales of Apple computers took off like a rocket. The desktop PC as a business tool was born (Source: Texas Instruments, 1981).

In 1981, Texas Instruments, at its Corporate Research Laboratory in Dallas, was developing a word processor program for the TI 99-4 PC that would soon be sold at a street price of $399. Business word processors from Xerox, IBM, Philips, and Wang were replacing the typewriter at a significant premium—costing $6,000 to $15,000 per workstation. The prospect of a less-than-$1,000 word processor (including the computer, necessary peripherals, and software) was perceived by marketing folks as the next "killer app" for PCs. Alas, the project was not completed (the 99-4 had some fatal flaws that caused it to crash in the middle of various tasks). Word processing on PCs did later became the killer app of the IBM PC and its clones. Why spend $10,000 for a word processor when you could achieve the same goal with a $3,000 IBM PC or clone running Word Star (an early disk operating system [DOS] version of word-processing software). Small, medium, and large businesses by the thousands bought IBM PCs and clones (most notably Compaq) because they could run Word Star and Lotus 123 (the next-generation spreadsheet program that replaced VisiCalc).

Killer applications put technology to work. The Internet became a crucial business tool with e-mail. Other important applications, such as the ability to do on-line research and e-commerce, have propelled the Internet revolution. *Figure 1* shows a few of the important milestones that have propelled the Internet to where it is today. The creation of the Mosaic browser followed by Netscape certainly qualifies as a killer application. Mobility may provide the next great boost to the growth of the Internet.

The promise of high-speed mobile Internet connectivity with 2.5 and 3G wireless applications provides some interesting food for thought as to how the Internet, liberated by wireless mobility, can change the way we communicate, do business, are entertained, or become educated. Before these sweeping changes can take place, various regulatory and security problems remain to be solved.

Restraints to E-Commerce Transactions

Governments and the private sector in countries around the world are grappling with issues that could slow down or even kill the growth of the wired and wireless Internet. Issues such as e-commerce taxation, privacy, and transactional security are at the heart of most of the trade-related discussions domestically and internationally.

Taxation

Taxation is a sticky problem. Nations, states, cities, and local communities are very concerned that their funding, which is based, at least in part, on sales tax revenue, will be eroded by e-commerce, which is protected under a Congressional moratorium on taxation in the United States. European Union member states, which rely on value-added tax for a large proportion of their budgets, have proposed that value-added tax (VAT) be imposed upon companies that conduct business over the Internet in their country. Such an imposition could cause a U.S. e-business to simply refuse to sell products to anyone residing in any of the 15 E.U. member states. This would clearly be a consequence that would disrupt the rapid growth of e-commerce and the growth of the Internet in general.

In *The Economics of Electronic Commerce*, Choi et al states that "to at least maintain the current level of tax revenues, state and local governments need to figure out how to apply existing rules of taxation to electronic transactions…similar situations are found globally." The Texas Sales and Use Tax Definition Sec. 151.005 defines tangible personal property (subject to taxation) as "personal property that can be seen, weighed, measured, felt, or touched or that is perceptible to the senses in any other manner." From this definition it appears that packet data is tax exempt. In another example, a purchase made out of one's state of residence is normally tax-exempt. What then is to prevent an Internet retailer from locating in two or more states and always doing "out-of-state" business? Clearly, tax laws will have to be changed and e-commerce defined before it can effectively be taxed.

Privacy and the Safe-Harbor Agreement

Privacy has become a hot issue, particularly in trans-Atlantic trade. European privacy rules allow individual member states to set their own levels of privacy protection for their citizens. This is a significant potential problem for U.S. e-business companies doing business with European customers. If a U.S. e-business company does not protect consumer data according to the laws of the member state in which their customer resides, they could be in violation of the law.

FIGURE 1

1989 HTML Invented; 1991 NFS Allows Commercial Use; 1993 Mosaic Browser; 1994 Netscape Browser Released

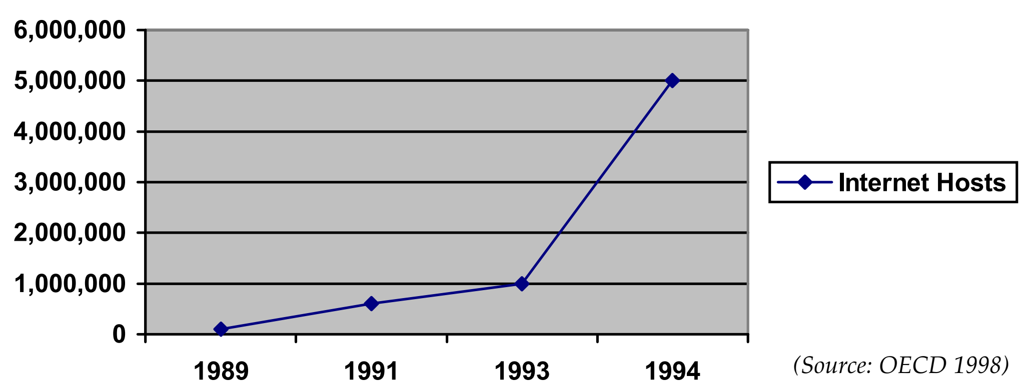

(Source: OECD 1998)

The U.S. Department of Commerce and the U.S. Trade Representative have successfully negotiated the issue with the European Commission and the member states to establish a "Safe Harbor" agreement. U.S. companies wanting to do e-business in Europe can sign on to the Safe Harbor agreement and then follow one uniform set of privacy requirements.

Network Security

Choi et al states that "unsecured transmission on the Internet is often sited as the main deterrent (to) rapid growth of electronic commerce." Add mobility to the equation, and the problem grows. The government tends to get involved at both ends of the security spectrum. On one hand, in the name of law enforcement and national security, there have been restrictions in place on the export of products with advanced encryption (a necessary element of security). Fortunately, the U.S. Department of Commerce and other agencies involved have realized the link between restricting encryption algorithms and the disadvantage it places on U.S. industry, not to mention the potentially harmful effects on the economy caused by inadequate levels of security.

At the other end of the security spectrum, the Department of Commerce's National Institute of Standards and Technology (NIST) and the National Security Agency (NSA) are participating, along with other governments and industry, to develop a "Common Criteria" for network security. The Common Criteria document states that "the Common Criteria represents the outcome of efforts to develop criteria for evaluation of IT security that are widely useful within the international community...and opens the way to mutual recognition of evaluation results." One of the goals of industry is to develop the Common Criteria as a consensus standard accepted by all parties and not regulated by national governments.

The Evolution of New Applications

As with the early computers, wireless applications will evolve together with enabling technology. An article in *TelOSSource* dated March 2001, Susan Weese said that "the wireless industry is standing astride another technology migration chasm. The market drivers for faster, content-enriched wireless services require a leap from existing networks and services to an entirely new business model. This transition is far more complex than the jump from analog to digital services just a few years ago. Delivering new services and enabling this wireless convergence between telecommunications, information technology, media and content will be a difficult and potentially rewarding task for enterprises…along with (managing) ordering and billing of new services, content providers must develop high quality, multimedia content to feed those new services "

Steven Hayes of Ericsson, who chairs a subcommittee in the Third Generation Partnership Project (3GPP), which is made up of wireless manufacturers and carriers, said, "I think one of the key issues with transition to the wireless Web will be the shift from a telco model to an Internet model of services. In the past, telecommunications services have been referred to as a 'walled garden,' and within that garden, the subscriber has good services available, but only services a specific telecommunications company provides. With the transition to the mobile Internet, there will be greater opportunity for third party service providers." *Figure 2* shows forecasted numbers of subscribers for three different 3G services (Source: *TelOSSource*, March 2001).

Japan's i-Mode Business Model

Some of the multimedia applications developed for 2G systems that will likely transfer to 3G systems can be found today in NTT DoCoMo's i-Mode service. The always-on Internet service over cell phones has grown from 6 million subscribers to 20 million in one year. Applications have developed rapidly, providing a good look at what people will pay for. One of the surprises has been a service that provides a "Pokemon" screen saver to cell phones. Literally millions of subscribers pay a little more than U.S.$2.00 per month to download updated Pokemon characters.

Wireless Week's Brad Smith (March 12, 2001) said, "To understand i-Mode's acceptance it is necessary to look at the serv-

FIGURE 2

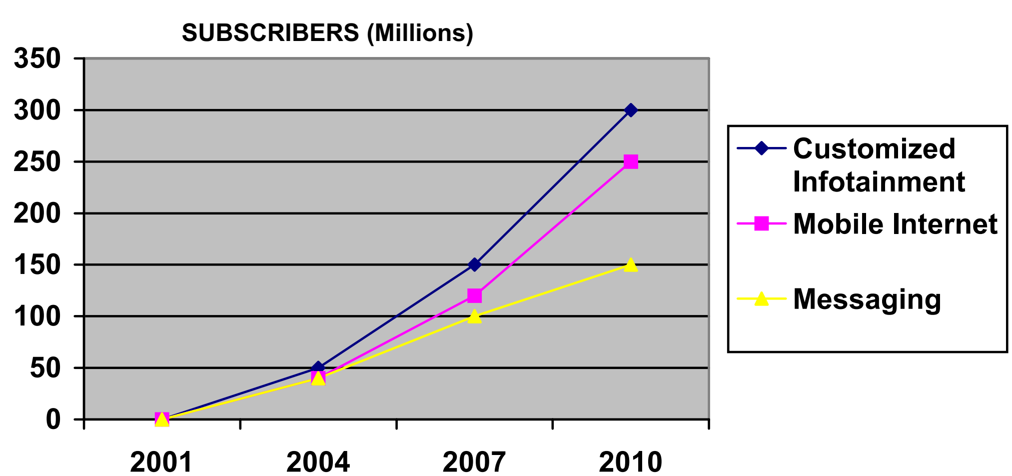

ice both as a technology and a business model. i-Mode is packet-based unlike the circuit-switched solutions that most Wireless Applications Protocol (WAP) carriers traditionally offer….What i-Mode offers to end users is rich content. DoCoMo has applications partnerships with more than 795 companies and 39,150 Web sites." There is no charge for Internet access, instead Web sites can charge for access. Some sites are free, but those providing value-rich content—such as news services, maps, and games—charge as much as U.S.$3.00 per month.

DoCoMo is taking i-Mode around the world with investments in AT&T Wireless (U.S.$9.8 million), KPN Mobile in the Netherlands, and Telecom Italia Mobile. DoCoMo is also leading everyone in the deployment of 3G services. They are building a wideband code division multiple access (WCDMA) network in Japan that will begin trials in May and be ready for commercial service in October this year—well ahead of any other carrier.

In a *Wall Street Journal* article (December 11, 2000), David Pringle said, "i-Mode is a huge hit in Japan, now the big question: will it travel well?…translating i-Mode into other languages is probably the least of DoCoMo's worries. Analysts believe that a lack of compatible handsets and suitable networks will make it difficult for DoCoMo to replicate the service in Europe…(But) The company has a gold mine of information about i-Mode that could be used to design a more compelling wireless Internet experience in Europe."

Killer Apps

So, what will make 3G phones become, in the words of the International Telecommunication Union [ITU], "an indispensable life tool carried by everyone, everywhere, just like a wallet or purse"? For starters, what if you didn't need a wallet—or purse, or cash, or credit cards, for that matter—to pay for things? The electronic wallet is available today in, of all places, Estonia. Drivers there can use their Global System for Mobile Communications (GSM) phones to pay for things such as parking fees.

The success of the electronic wallet depends heavily on network security. It also will require re-thinking of how payment transactions are handled. Once these issues are worked out, electronic payment on 3G phones, perhaps by communicating with cash registers through a Bluetooth connection, may well be the first "killer app".

Christopher Wu of Yahoo!, in a December 11, 2000 *Wall Street Journal* article, said, "The best way to get consumers on the mobile Web is to help them use it for old-fashioned shopping, in-person, at stores…..E-wallets are a necessity for any type of m-commerce. There's no way you're typing in all that (credit-card information) every time you want to buy something." With a sufficient level of security protection, the e-wallet will certainly become the basic tool for m-commerce.

Music via Wireless

Shawn Young, in a *Wall Street Journal* article (December 11, 2000), said, "Music is a strong candidate to become a killer app, analysts say. Mobile devices could combine the enormously popular success of online music with the portability of a Sony Walkman." Ericsson currently makes a Moving Pictures Experts Group (MPEG) Layer-3 (MP3) music module for GSM phones. The added bandwidth of 3G networks (384 kilobits per second [kbps] to 2 megabits per second

FIGURE 3

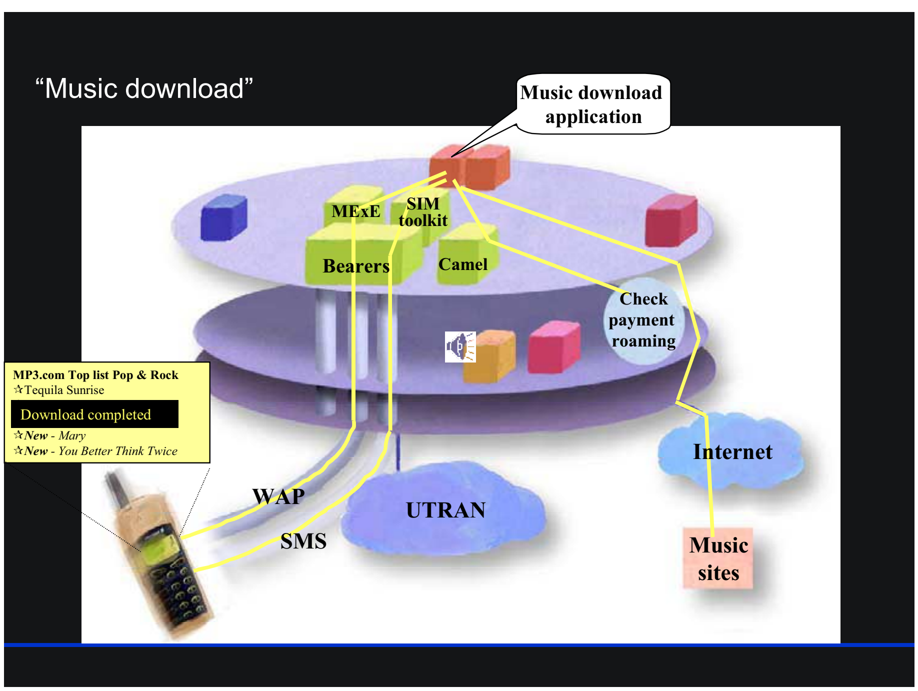

[Mbps]) could indeed make mobile music a killer app. *Figure 3* shows how a music purchase would be transacted over a mobile network.

Location-Based Services

The Federal communications Commission (FCC) has ruled that wireless carriers must provide location information for 911 calls. The ruling calls for ultimately locating a mobile phone within 50 to 150 meters. Today, carriers can identify which cell site the call is coming from, but that guarantees only that the call is within one to several miles from the base station. That information would not help an ambulance driver find someone in need of emergency medical treatment. There are essentially three technologies being proposed to provide pinpoint location information. The first is network-based and uses triangulation based on the time it takes for a signal to travel to two or three base stations or the time to reflect from known objects near a base station. The second technology is handset-based, using the Global Positioning System (GPS) in the handset. The third is a hybrid technology that requires modifications of both the handsets and base stations. The cost of installing any of these technologies is estimated be in excess of U.S.$1 billion.

Leave it to entrepreneurs to find a commercial opportunity in an FCC mandate. Young, in the *Wall Street Journal*, said, "Analysts expect a proliferation of location-based services that tailor the content to a user's location, providing directions to the nearest theater, say, or the neighborhood Italian restaurant or service station. For such convenience, the theory goes, tightfisted Web browsers might finally be willing to fork over some money."

Games and Other Revenue-Generating Applications

NTT DoCoMo reports that more than half of the i-Mode pages accessed by users are cartoons, games, horoscopes, and other entertainment services. Considering that many young people use computers for playing games, it is easy to imagine that entertainment, including games, is likely to contribute to more than a few killer apps.

In his *Wall Street Journal* article, Young identifies several potential killer apps, such as messaging, which would provide a natural extension of wired e-mail. In Europe today, there are 15 billion messages sent per month on GSM short messaging services (SMS).

Young mentions other applications, such as comparison shopping, movies (downloaded video), and a very important area of "Business Missions." He says, "Many experts believe the real money lies in business use of the mobile Web. CreSenda Inc., a closely held Los Angeles company is beaming real-estate listings, maps and images to real-estate agents….More broadly, industry experts say, giving people mobile access to corporate networks—so they can handle normal business transactions from lightweight, easy-to-use devices—could be a real killer application on the wireless Web."

Revenue generation for the carriers and the content providers will probably look much like the i-Mode model of micro-payments. DoCoMo adds charges for Web site access to the user's monthly bill, takes a nine percent commission and sends the rest to the content providers. *Figure 4* shows how 3G revenue may be earned by carriers and service providers.

FIGURE 4

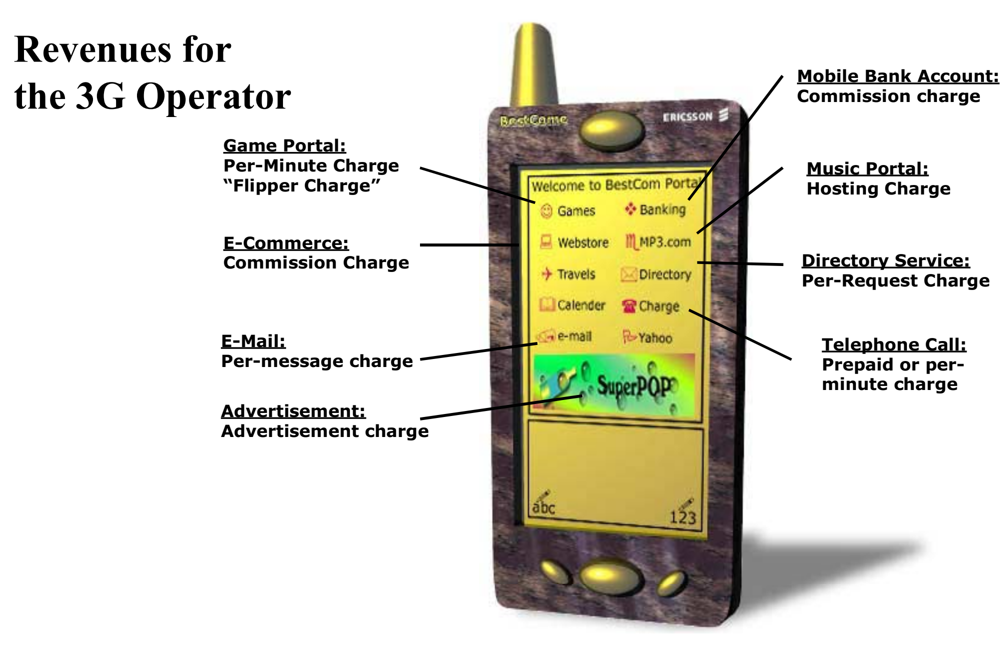

Conclusion

That 3G wireless services will change the way we live, do business, and are entertained is easy to imagine. The mobile Internet will also change the way we are taxed and will pose vexing problems for regulatory authorities. Getting there is full of obstacles that must be overcome. Operators, equipment manufacturers, and content providers will need to manage new challenges such as the cost of 3G licensing (U.S.$230 billion in Europe), the anticipated cost of network build-out, and the cost of developing content that will eventually produce the revenue to reach profitability. The transition from first-generation analog networks to second-generation (2G) digital services offers a model, but the transition to 3G and the mobile Internet is much more complex. Instead of simply introducing a better system for voice calls and later SMSs, as was the case with 2G, we are adding the Internet and mobile commerce. There will be mistakes and failures along the road to 3G, but as Columbia University Professor Eli Noam said, "a lot of people tried to fly before the Wright brothers came along."

The Evolution of Mobile Network Backhaul Architectures

High-Capacity Multiservice Networking for Mobile Backhaul Applications

Larry Keith
General Manager
Millimetrix Broadband Networks, Inc.

Introduction

The demand for new data services combined with the growth of mobile networks requires mobile operators to begin planning the evolution of their backhaul networks to support new network requirements. Additionally, the public's increasing reliance on mobile networks creates the need to improve network reliability and availability, while economic realities drive the need for lower-cost network solutions. These often-conflicting demands make it necessary for the operator to rethink the backhaul architecture and begin the implementation of scaleable, reliable, cost-effective backhaul network architectures.

This white paper examines the current state of mobile network backhaul architecture in both access and core networks, then looks at factors driving the network evolution toward higher capacities, new protocols, and improved reliability.

Current Mobile Network Architectures

Mobile network backhaul architectures are usually categorized into two segments, each with different network requirements. The *access network* segment includes the subscriber-to-base transceiver station (BTS) connections (the mobile wireless connections) as well as the connections between BTSs and the base-station controller (BSC). This part of the network is characterized by low bandwidth requirements, many distributed locations, and varying traffic demands.

The *core network* segment includes BSC–to–mobile switching center (MSC) connections and MSC–to–MSC connections. This portion may also include connections from BTS hubbing locations (subsequently defined) to the BSC and connections from MSCs to the wired telephone network. This network portion is characterized by high bandwidth requirements and fairly predictable traffic demands due to the concentration of traffic, and by the need for high reliability because so many BTSs are aggregated into a BSC.

Figure 1 shows the architecture for typical, second-generation (2G) mobile backhaul networks. Each of these network segments has unique requirements that drive the current mobile network architecture.

Access Network

The access part of the network includes subscriber-to–BTS (mobile wireless) connections as well as BTS–to–BSC connections; only the BTS–to–BSC connection is considered part of the backhaul network architecture. For 2G mobile networks, each BTS typically has one to four T1 or E1 circuits to the BSC; the exact number depends on the expected peak load of that BTS and the number of sectors. The communication in the access network is strictly hierarchical; the BTSs communicate only with the BSC, not with other BTSs. Each BSC may serve from a few to several dozen BTSs.

The current backhaul architecture in the access network uses either fixed wireless T1/E1 systems or copper-based T1/E1 circuits in either a point-to-point configuration between each BTS and its BSC (the "star topology") or a ring of T1/E1 radios or circuits connecting the BTSs to the BSC. Fixed wireless is used when line-of-sight is available, and the distance provides sufficient availability (typically 99.95%) for the service. Otherwise, copper facilities from the incumbent operator are used.

Core Network

The core part of the network includes BSC–to–MSC and MSC–to–MSC connections. In some networks, where it makes economical sense to aggregate traffic from several BTSs before connecting to the BSC, the connections from this "hubbing" BTS to its BSC may be considered part of the core network as well. In 2G networks, each BTS typically requires a digital signal (DS)–3/E3 connection to the MSC, and hub BTSs typically have eight to 16 T1/E1s or a DS–3/E3 connection to their BSC.

In the core network, communications is strictly hierarchical from hub BTSs to BSCs to MSCs. If the network has multiple MSCs, these operate in a peer-to-peer fashion, with each MSC switching traffic locally and sending traffic to the other MSC if required based on call destination. MSC–to–MSC connections are typically optical carrier (OC)–3/synchro-

FIGURE 1

Current Network Architecture

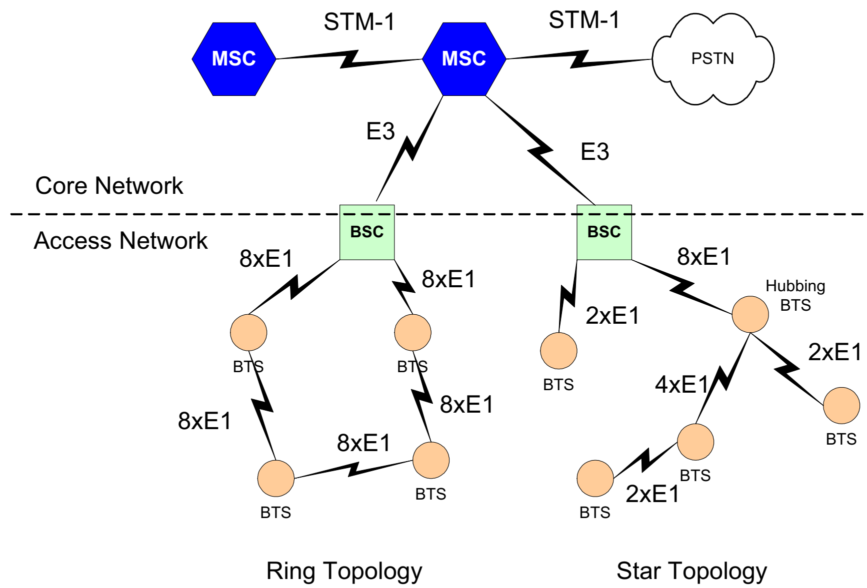

nous transfer mode (STM)–1. A mobile network may have anywhere from two to dozens of BSCs, and one to four or more MSCs.

The current backhaul architecture for the core network again uses a combination of fixed wireless and wired facilities. Due to the higher capacities required, wired facilities must be fiber-based, which may not be available outside of the urban core. Outside of urban core areas, fixed wireless is the best solution for the backhaul network. Because of the reliability needed for aggregated traffic, protected (fault tolerant) configurations are typically used for these facilities, either synchronous optical network/synchronous digital hierarchy (SONET/SDH) rings for fiber-based facilities or protected (hot-standby) configurations for fixed wireless facilities.

It should be noted that in both access and core networks for 2G systems, the traffic is strictly time division multiplexed (TDM); no packet-based protocols are used.

Evolution of Mobile Networks

Like much of the telecommunications industry, mobile communications technologies are constantly evolving. The increasing penetration of and dependence on mobile wireless communication necessitates better network planning. The advanced data services made possible through 2.5G and third-generation (3G) offerings will provide users with high-speed access to Internet and intranet applications from their mobile devices. Furthermore, operators and regulators are investigating how they can share infrastructure to reduce 2.5G/3G network deployment costs. These trends are described in the following section.

Continued Growth

The worldwide wireless subscriber base continues to grow, as shown in *Figure 2*. In the United States, which is behind much of Europe in mobile subscriber penetration, the September 11, 2001 terrorist attacks have accelerated subscriber growth, since individuals have realized the value of cell phones for security and personal use.

Studies have shown that the saturation point for mobile subscriber penetration is 70 to 80 percent in developed countries, primarily due to age and poverty. If this is true, there is still room for significant growth in many countries. This incremental growth will only increase the amount of traffic carried on the mobile backhaul networks.

Improved Reliability

As the mobile subscriber base continues to grow, consumers and businesses that become more dependent on wireless service will also become less tolerant of dropped calls, outages, and busy signals. Consequently, operators must incorporate network redundancy in planning future networks.

2.5/3G Data Traffic

The principal driver now under way in mobile wireless penetration is the introduction of high-speed mobile data services. Known collectively as 2.5G and 3G technologies, these services include general packet radio service (GPRS), enhanced data rates for global evolution (EDGE), 1xRTT (Radio Telecommunications Technology), wideband code division multiple access (WCDMA), and CDMA2000. 2.5G systems are generally characterized as having peak data throughputs of 64 kilobits per second (kbps) to 384 kbps, while 3G systems can support data rates of up to 2 megabits per second (Mbps). Compared to voice traffic that ranges from 16 kbps to 64 kbps, data traffic could require a tremendous increase in the capacity of the backhaul infrastructure. In addition, data traffic is "bursty," meaning that the average throughput is much less than the peak throughput, while voice traffic is constant throughout the conversation. Therefore, migration from the current copper-based access

FIGURE 2

Mobile Subscriber Growth by Country

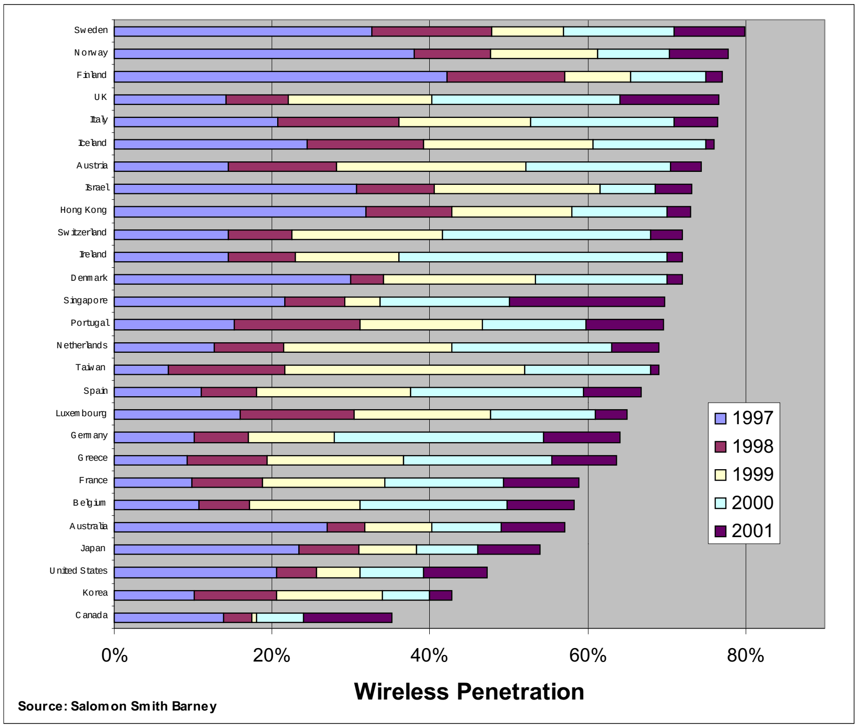

network infrastructure to 2.5G and 3G network will be a tremendous challenge.

To achieve the highest rates, users of 3G networks will need to be closer to the BTS than they need to be with 2G systems. This means that eventually the density of base stations must increase, creating a larger concentration of BTSs for each BSC (called a radio network controller [RNC] in WCDMA). This higher density will increase the traffic in the core network from each BSC/RNC.

Furthermore, while the current mobile network is built around TDM technologies, 2.5G/3G will introduce packet-switched communications into the backhaul network. Packet switching is a more effective method of supporting bursty traffic, as statistical multiplexing takes advantage of the "lulls" in data conversations to support other users, allowing more users to be supported with the same bandwidth. Asynchronous transfer mode (ATM) is the packet switching protocol currently used in 2.5G and 3G systems, although Internet protocol (IP) is widely expected to replace ATM in future systems, gradually moving from the core network into the access network.

The plesiosynchronous digital hierarchy (PDH) and SDH transport architectures currently in use are optimized for TDM traffic and do not efficiently support packet-switched communications. Any future network architecture must include efficient support for packet switching to realize the economic benefits of a shared (voice and data) transport network architecture.

Though 2.5G/3G networks will become the infrastructure of the future, they will be built on top of today's 2G networks. The migration will give subscribers ample time to upgrade their handsets, which will increase the demand for backhaul from each BTS site.

Infrastructure Sharing

Return on invested capital is of great concern for mobile network operators as they plan the rollout of 3G networks. Having paid so much for licenses, many operators cannot ever hope to earn enough return on their investment without some regulatory relief. While such relief may take the form of reduced license fees or tax reductions, many countries are considering "infrastructure sharing" to be the best form of relief, since it also reduces the impact that new networks would have on the environment. Infrastructure sharing would let operators share towers, antennas, base stations, sites, and access, reducing the cost of deploying new networks.

To the extent that regulators and carriers adopt infrastructure sharing, the need for higher-capacity backhaul will be amplified. With infrastructure sharing, two to as many as five or six operators could potentially co-locate on the same tower, creating the need for two to six times more bandwidth in the shared backhaul infrastructure.

Primarily in the United States, mobile operators have been selling their tower portfolios to specialized tower companies. This trend may expand to Europe and Asia as operators focus on their core business. Tower companies act as tower landlords, leasing space to multiple mobile operators.

Leasing tower space, such as the 3G regulatory relief previously described, serves to effectively amplify the backhaul bandwidth requirements at each tower.

Future Network Architecture Requirements

Given the evolution just described, it is clear that the next-generation backhaul network architecture must offer several improvements over today's network, while maintaining current aspects such as fast deployment, cost-effectiveness, and low delay for voice traffic. These required improvements are described in the following sections.

High Capacity
The future mobile network architecture must support higher capacities. BTS backhaul will require additional T1s, but capacities must increase dramatically in the core network. Whereas today DS–3/E3, and in some cases OC–3/STM–1, is sufficient, OC–12/STM–4 will become the base transport capacity. In the dense urban core, OC–48/STM–16 will be used.

Multiservice
The future mobile network architecture must provide efficient support for multiple services, including TDM as well as packet-based services. TDM will remain in the network to support current 2G networks, while 3G deployments will initially require ATM and, later, IP. While packet over TDM is an option, many operators will want the economies of native packet support in their transport network.

Multiple Transport Media Options
The move away from copper-speed to fiber-speed networking creates enormous challenges for operators. Fiber does not exist in the metro area outside of the urban core, so other alternatives must be considered. Broadband wireless has been and will continue to be an excellent alternative, and emerging free space optics solutions can also have a role for short distances. The challenge for the operator will be to integrate these multiple transport media options into a seamless network architecture.

High Availability
To provide the highly reliable network that customers expect, operators will need to plan around the inherent limitations of the transport media. For fiber, this means alternate routing—where this is not feasible, fiber combined with broadband wireless or free space optics. For broadband wireless links, protection switching must be considered, and outages due to rain fade must also be planned for. Advanced radios use adaptive modulation to reduce the data rate while still maintaining the connection during the worst rainstorms. For free space optics links, fog and snow fade issues can be solved with adaptive coding, again to reduce the data rate while maintaining a reliable connection.

Flexible Configurations
A high-availability network ensures that each link is almost always available by providing multiple paths for each connection. Ring or mesh architectures provide multiple routes for customer traffic, ensuring that the network stays up even when entire sites may be lost due to accidents, power outages, or natural disasters. Mesh architectures are especially flexible, allowing the operator to grow the network based on traffic demand rather than attempting to pre-plan the entire network, which can be the case with ring architectures.

Fast Deployment
The frenetic pace of network deployment will only intensify as mobile wireless networking continues to evolve. The slow pace of fiber deployment will not support the market demand. The future network architecture must use innovative solutions such as broadband wireless and free space optics to quickly deploy and make new sites operational.

Quality of Service (QoS)
With all of the changes required of the mobile backhaul network, delay requirements must not be forgotten. The future mobile network architecture must provide the strict QoS required for voice conversations, whether carried as TDM or as packets. Data is more tolerant of delay and packet loss, but the operator must still be able to provide service-level agreements (SLAs) for data to support corporate applications.

Low Cost
Finally, all of this new functionality must be economical. Going forward, operators must be even more cautious with their money for at least two reasons: operators already spent a tremendous amount on licenses, and operators don't know how much consumers will be willing to pay for mobile data. Balancing capital expenditure (CAPEX) with revenue potential has become more critical. A reasonable objective in the core network would be to upgrade from OC–3/STM–1 to OC–12/STM–4, a four-times bandwidth increase, with no more than a two-times increase in CAPEX and recurring costs.

Future Mobile Network Architecture

One concept of a future mobile network architecture is shown in *Figure 3*.

This architecture contrasts with the current architecture shown in *Figure 1* in several important ways. The future network architecture uses a variety of transport media at STM–4 (622 Mbps) rates in the core network, as opposed to the STM–1 and E3 rates of the current architecture. The access network uses hubbing BTSs to a greater extent and uses STM–1 links in a ring topology to connect hubbing BTSs to the BSCs for higher capacity and higher reliability.

Unity‘ for Mobile Network Backhaul

One solution that meets all of these requirements is the Unity‘ system from Millimetrix Broadband Networks. Unity is a multiservice, multimedium transport system that is ideal for mobile network backhaul. It facilitates the buildout of hybrid, high-capacity, optical-rate networks that combine fiber or "fiberless" links, such as licensed radio channels and unlicensed free space optics links, to create a unified transport network. The Millimetrix transport layer incorporates proprietary Optilinx‘ technology that maximizes the reliability, dynamic nature, availability, and constrained capacities of nonfiber media. Millimetrix's Optilinx technology provides a multiservice, multimedium platform that supports both legacy and emerging services and creates the modularity and network scalability that give the carrier the competitive advantage needed in a dynamic market.

FIGURE 3

Future Network Architecture

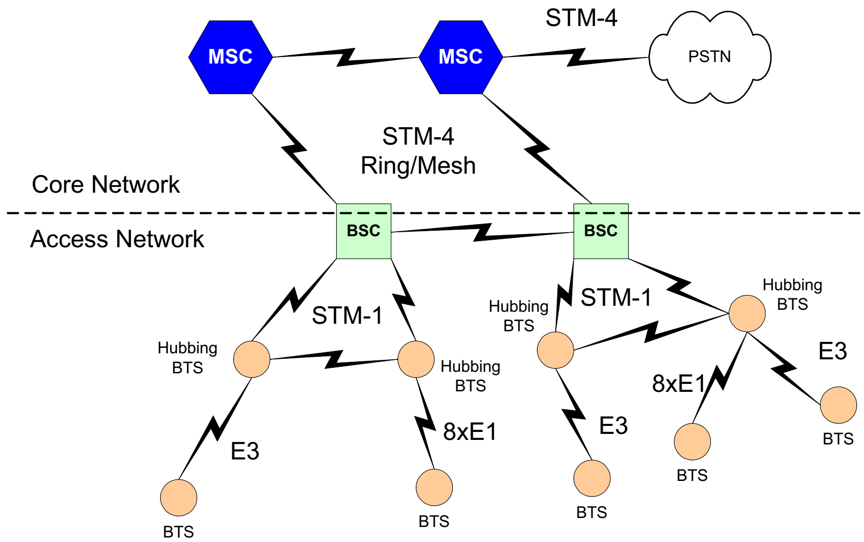

The Unity architecture is comprised of two network elements: an indoor unit (IDU) and an outdoor unit (ODU). The IDU includes service interfaces, service provisioning and monitoring, service bandwidth management, multiplexing and switching, and the Unity network-management interfaces. The IDU has a modular architecture and supports OC–3, OC–12, Fast Ethernet, and Gigabit Ethernet. The multiplexed and switched traffic is then transmitted to other Unity sites using existing fiber or a “fiberless” alternative, provided by Unity ODUs. Millimetrix offers two types of ODUs: a radio frequency (RF) ODU using a single licensed radio channel (50/55/56 MHz), and an optical ODU using an unlicensed free space optics channel. Each Unity ODU type supports a maximum capacity higher than 622 Mbps.

A service provider deploying a next-generation network with Unity will benefit from being able to exploit available fiber along with the other alternatives provided by Unity and from being able to scale the network as customers are acquired. Unity enables the immediate generation of service revenue while giving the carrier reduced CAPEX per bit of traffic.

Conclusion

The mobile network backhaul architecture of the future will be very different from today’s network. In addition to supporting more voice traffic and efficiently supporting new high-speed data services, the network will need to offer improved reliability at a lower relative cost. This network will take advantage of fiber but not be limited by it, using broadband wireless and free space optics wherever it benefits the operator. This new mobile backhaul network will be a key enabler of profitable services for the forward-thinking operator.

Content Revenue Solution for Mobile Internet Services

André Kopostynski
Director of Marketing, Customer-Management Solutions Services Unit
Ericsson

Overview

The market for wireless Internet applications is set for explosive growth both in the United States and around the world. The Yankee Group predicts that by 2005, 56 million people or 20 percent of the U.S. population will be regular users of the wireless Internet over a voice-enabled device representing 15 percent of carriers' revenue[1]. Moreover, Merrill Lynch estimates by the end of 2003, there will be approximately 47 million mobile Internet subscribers representing 25 percent penetration of U.S. wireless subscribers[2]. Service providers wanting to capitalize on the demand for wireless Internet will need to invest large amounts of capital in network expansion and licenses as new technologies such as third generation (3G) requires massive replacement in infrastructure as well as new radio spectrum.

This white paper discusses some of the major challenges that wireless operators and service providers are facing as they transition themselves from mobile telephony providers to next-generation carriers with service offerings that combine voice and data, content, and other value-added applications to an increasingly sophisticated and demanding user base. New network technologies spell new business opportunities. This paper will scrutinize how these opportunities are challenging the existing operations and business support systems (BSSs) in which carriers have invested. Additionally, the carrier support system features that are required to support the emerging value chain, where new revenue streams will be made up of several types of pricing schemes, such as airtime or data transfer fees, content fees, transaction fees for m-commerce and advertising, and other sponsored revenues, will be featured.

To justify the significant capital investments that 3G requires, operators must seek ways to maximize the revenue potential of their data offerings. After remarkable growth of the traditional voice subscriber base, where some countries have reached penetration rates in excess of 75 percent, wireless carriers now are witnessing declining average revenue per user (ARPU), increasingly saturated markets and segments, intense competition, and very few means, if any, to differentiate their offerings from rival carriers.

As virtually every competitive communications service provider is struggling to justify their network investments, it has become apparent that flat-rate, all-you-can-eat pricing schemes for accessing data networks and the Internet are not necessarily sustainable business models. Differentiation in the marketplace has been hard to achieve on any factor other than price. The downturn in the dot-com and broadband industries proved that "market-share gains at any cost" or "if you build it, they will come" were not sound business strategies, either. New entrants and intensified price competition made profitability virtually impossible to achieve, at least within the expected time frames set by venture capitalists. So given this history lesson, wireless carriers must ensure that they don't become just "dumb" transmission pipes for content and application services because, as recent times have shown, the only means for differentiation will be price. With ARPUs in mind, carriers now are re-evaluating their business strategies and market positions to capitalize on the explosive growth in demand for wireless Internet and content services projected by industry analysts.

Protect Existing System Investments

Wireless carriers have traditionally viewed themselves as providers of mobile voice services—an extension or complement—to traditional wireline telephony service. Only recently, due to economies of scale and the maturing of wireless technologies, has mobile telephony been regarded as a viable substitute to its wireline counterpart, at least from an end user's perspective. So naturally, the focus has been on customer acquisitions and the build-out of circuit-switched network capacity to handle more voice traffic. This focus is now all about to change.

Now, operators must carefully re-evaluate their business strategies and market positions to capitalize on the explosive demand for wireless Internet and content services. As service providers transition themselves from the second generation (2G) by means of 2.5G, to 3G broadband wireless data capabilities, they must redefine internal operational and business processes. This self-examination is essential to accommodate new requirements resulting from the dynamic Internet protocol (IP) world and growing customer expectations.

One of the greatest challenges facing these carriers is centered on preserving their existing operations support system (OSS)/BSS environment. Operators have invested substantial capital in these systems to automate and streamline operations and businesses processes, as most operators are not inclined to invest more capital in replacing their legacy back-office systems. A billing and customer-care system typically accounts for 20 to 30 percent of total OSS and BSS investments and often requires a significant number of dedicated systems specialists to run and maintain this complex environment.

Although support systems have evolved over time to accommodate newer services, such as CallerID®, voice-mail, and others enabled by 2G digital networks, their core functionality has been to support millions of customers with circuit-switched voice networks. Voice services have generally been priced on three simple dimensions: duration, destination, and time-of-day, where duration is the "real" metric for network capacity utilization. With the introduction of packet-based networks and services, these pricing metrics are no longer adequate.

Given the enormous capital outlay required to transition into 3G, carriers will need tools to effectively price new packet-based services based on nonstandard dimensions such as customer's perceived value, competitive forces, and network utilization. To further complicate things, carriers will also need to include partners in its value chain—adding the requirement for relationship management. Traditional legacy systems were simply not intended to support, or capable of supporting, this new industry paradigm.

So what are the viable alternatives? A total replacement is certainly an alternative; however, the sizeable capital outlay will place enormous stress on the finances of the carrier—and it's risky, given the uncertainties of future system's functional requirements.

Another approach is to invest in a dedicated IP solution and run it in parallel with the existing legacy system. This certainly will reduce the risk and the cost, but there are some apparent drawbacks, such as operational inefficiencies of having to operate and maintain two complex systems as well as making it difficult to effectively bundle voice and data services into attractive, convergent offerings. Additionally, there is the possibility of having to generate two separate invoices to the same customer (one for data services and one for voice).

A more cost-effective and even less risky alternative may be to explore enhancement opportunities to the existing system environment. Because many key business processes will essentially remain the same (invoicing, payment handling, taxation, account reconciliation, etc.), why replace the entire environment? The core changes in transitioning to 3G will stem from the network side and how to collect, process, and rate the usage records, price the delivered service, and share revenue where partners are involved—before the information is passed on to the billing and settlement systems.

Existing OSS/BSS systems support the voice business well, which will also remain the key revenue source for some time to come. However, with the introduction of IP networks and content-based services, carriers need a dedicated focus on content mediation.

Content mediation is a term applied to represent the collection, aggregation, and processing of raw usage data generated in IP–based networks such as the Internet. In a traditional circuit-switched voice network, mediation is relatively straightforward task, because the central point of collection is the switch itself, which typically generates all the required information to bill for the call; this information is often referred to as CDRs or call detail records. However, in an IP network, this represents a much bigger challenge, because there is no single device from which to collect all the necessary usage details. To capture meaningful usage information, IP mediation platforms must collect from several devices and elements in the network. And since there is no true industry-wide standard with regards to the information such IP elements should generate, hardware vendors have taken a proprietary approach to this issue. Fortunately, the IP industry recognized this problem early on and formed several standardization bodies that are now making great progress and are expected to have a solid framework in place in the near future. Another challenge has been to design mediation solutions with capabilities to collect and process millions, if not billions, of daily IP events. This has been a major impediment to full-scale adoption by the industry, but we now are overcoming these hurdles and actually proving that IP events can be collected, aggregated, and processed even at extremely high volumes.

Raw usage data is usually collected directly from the elements themselves, such as IP switches, gateways, routers, and application or Web servers. A content-mediation system should collect this information in real time, aggregate and filter out unnecessary and redundant data, and then process the usage information and prepare it for other systems, such as rating and billing, network management, fraud-detection applications, etc. It should also have the capability to automatically distribute the usage data to comprehensive network-mediation system for enhanced processing, e.g., apply additional business logic so that cross product and/or volume discounts can be calculated. Content-mediation systems are unique systems that must to support a number of network elements, such as the following:

- AAA servers running RADIUS or TACACS+
- Web servers
- Mail servers
- Firewalls
- LDAP servers
- DNS
- Routers/IP switches (SNMP, NETFLOW™, etc.)
- Gatekeeper/gateways for WAP/SGSN/GGSN

The collection method should support user datagram protocol (UDP), transmission control protocol (TCP)/IP, file transfer protocol (FTP), software agents, simple network management protocol (SNMP), and remote authentication dial-in user service (RADIUS)/TACACS+. Software agents must enable the real-time collection and data-record creation of billing information from server application log files. This is applicable for games, news applications, and other services that have user-defined value.

To support the services enabled by faster data transmissions, there is a need to collect data from many sources and handle large data volumes, whether from the network elements or from the content elements—or a combination of the two—such as Web and application servers. To assure revenue streams in this complex flow, mediation-processing capabilities are extremely critical. If the switches send the same IP detail record (IPDR) twice to the mediation system, this needs to be identified. The mediation system needs to provide an easy-to-use interface to ensure that faulty IPDRs in a file can be detected and corrected before being forwarded to the billing system.

Formatting is important when collecting IPDR files from a variety of network elements (switches, access nodes, gateways, servers, etc.). This is to protect the billing application from too much change so that operators can quickly bill for new services using their existing systems.

Different formats are collected, converted into a standard file format, and thereafter distributed to the downstream systems. Validation capabilities are also important to secure the new revenue streams. The mediation system must enable the operator to do the following:

1. Validate that the file contains billable data records
2. Consolidate partial records; filter and aggregate data records
3. Stream/burst or filter data records to downstream systems

When collecting information from routers and other Internet elements, it is vital to reduce the superfluous information in order to process only the information that is relevant for rating and billing purposes. When collecting router data, for example, the filtering should be performed within the content-mediation server. If a software probe or IP agent is deployed in a content or application server, the probe should filter out unnecessary information and send only distinct or relevant billing information to the mediation server.

The burst functionality enables the operator to send the same data records to several downstream systems. This feature optimizes the use of business intelligence applications, fraud-analysis software, as well as other support applications.

The content-mediation software should be configurable as either a standalone application or as a fully distributed solution. The key advantage of having this flexibility is that when huge amounts of data are collected in real time, specific nodes close to the network elements can perform some initial filtering and aggregation of data, which in turn can be sent to a central mediation server, which deals with the final aggregation and consolidation of all the billable data. This means that there are no real limitations as far as data and performance are concerned

Many carriers, especially the incumbents, have invested significant resources in their existing support systems, and naturally they seek to protect these investments whenever possible. As already argued, existing support systems such as billing and service provisioning are not agile and flexible enough to support the emerging business models of tomorrow, yet their core functionality is still needed to automate and carry out the processes that generate the bulk of the revenues. Making changes to or replacing such systems often requires substantial resources. Consequently, decisions to replace systems or implement any changes must be carefully scrutinized and produce an attractive return on investment (ROI) potential.

Accelerate Time to Market

Increased competition in the mobile voice service sector has resulted in a more homogenous, commodity-like offering. Carriers experience great difficulty in differentiating their service offerings in the marketplace. Essentially, pricing or carriers' "bundling of minutes," and to a certain extent, their networks' overage area, are the only means that they have to stand out from competition. New revenue sources are desperately needed, not only to rise above the competition, but also to justify the huge capital investment in infrastructure replacement and spectrum licensing. It is also clear that this "new" revenue must come from existing customers as well as new ones. Time-to-market and revenue

FIGURE 1

Filtering of IP Events

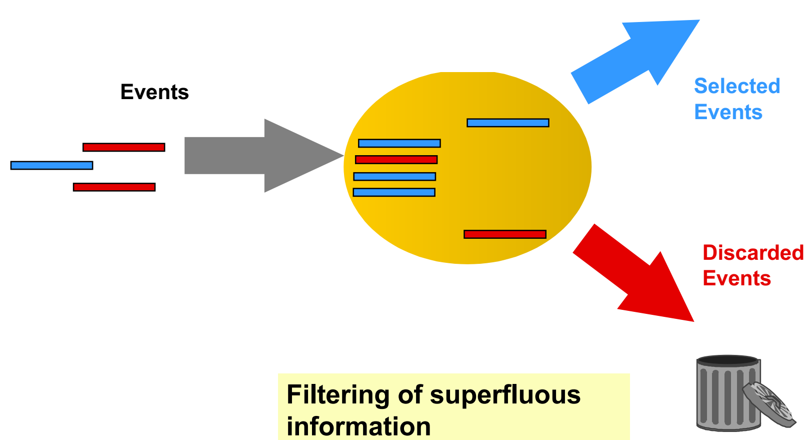

FIGURE 2

Overview of Ericsson Customer-Management Solutions

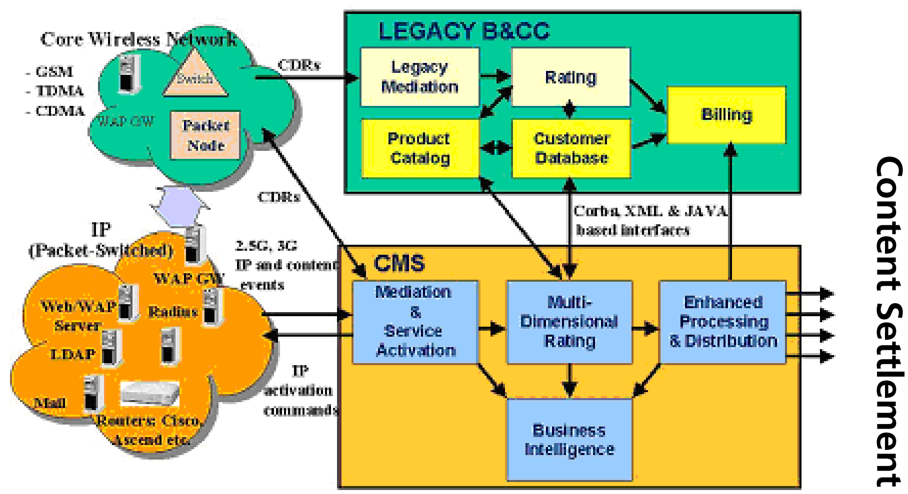

pressures are growing quickly and competition is intensifying even more.

The wireless industry has recognized that branding can often be a potential differentiator. Verizon and Cingular seek to distance themselves from the outdated "baby bell" image, using new and creative names to shed this out-of-fashion skin. North American wireless carriers have the added challenge of delivering mobile IP connectivity and content services to an already experienced Internet user base. Americans are familiar with the dynamic nature of the Internet, and their expectations will naturally be extended to the mobile Internet as well. However, it is extremely important that carriers effectively educate their customers and set realistic expectations as they introduce these new services. The user experience, at least in the near future, will be vastly different from desktop "surfing."

One potential key to the successful introduction of wireless mobile data service will greatly depend on the effective management of customer expectations. Carriers need to "flirt" with various business models and take proactive steps to become more "customer-intimate." Japan's DoCoMo® has witnessed tremendous success with its i-Mode® platform. Within 13 months, it attracted more than five million subscribers. Its success can be attributed in part to DoCoMo's willingness to experiment. Service providers often do not know what users want from wireless portals, and the only sure way to find out is by exploration.

Regardless of the chosen market entry strategy and target segments, wireless carriers need cost-effective and flexible management tools for introducing new IP–related services and products. They need to create and adopt an environment where business and marketing decisions can be implemented "on the fly." First-mover advantage can be extremely valuable in terms of quickly building a critical mass and a sizable market-share. Carriers can then expand the scope to attract new, targeted customers with highly personalized application and content services.

Network operators and service providers will require support systems that not only automate and carry out many of the operational processes, but also a system environment that can cost-effectively evolve with changing market needs and business drivers. Next-generation OSS/BSS software needs a layered architecture to enable swift and drastic changes in business models and service offerings. Such support systems need to be easily configurable and driven by rule-based logic. In addition, the need for intuitive user interfaces is critical to ensure effective administration of the support systems.

The amount of new IP–based services that general packet radio service (GPRS) and 1xRTT networks can provide is only limited to operators' creativity and the perceived value in the marketplace. Intense competition demands that new services should be introduced very rapidly—in days, rather than months. Customer agreements will also change, and different promotions may be active at the same time. The rating or pricing logic must be flexible enough to support all types of measurement units and pricing schemes—including duration-based, volume-based, subscriber-based, value-based, or any combination of these. Because of this, the rating architecture needs to handle any possible service and without additional programming. The volume of events created from IP–based services will be much higher than circuit-switched voice traffic; therefore, the performance and scalability of the rating core must reflect this and allow for massive scalability.

IP and content usage data records will be received in multiple formats, from multiple sources, including import of prerated data. The tariff calculation mechanism must be distributed to the content servers, thereby removing the burden of updating an ever-expanding product catalog and its associated price list in the central billing system. The price or the price group information must be placed in the usage data record.

Moreover, there is a need for a flexible, set-up driven interpretation of these records. Events must be rated in real time.

One single traffic event must be allowed to generate multiple charges in an environment where more than one party may be charged or compensated. For example:

- Revenue sharing with content partners and advertisers
- Charge for both volume and duration and the associated QoS

Many customers, especially business customers, are seeking special agreements with service providers because of their unique needs. There is a need for a straightforward pricing structure for the standard services offered, as well as unique pricing agreements that are discounted based on standard service offering. This will significantly reduce the maintenance cost of price lists and customer agreements.

Furthermore, services will cross-subsidize each other, and cross-discounts will be used to encourage customer loyalty. Mobile Internet access, personalized content, and value-added services need to be bundled into product packages to create "sticky" offerings. The package will also need to have a lower price than the sum of the integral components. Therefore, there must be a user-friendly approach to version handling of the price lists.

There is also a great need for service activation functionality when delivering IP services over wireless infrastructures. Services such as mobile virtual private network (VPN) with access to corporate Intranets and databases need to be activated rapidly and securely. To achieve operational efficiency, activation of services must be automated and executed in real time. It should not take more than a few seconds to activate a new service because customers expect a service to be available immediately after it has been ordered.

Activation is primarily about receiving activation requests from order-management systems, translating the requests into network-element commands and distributing these commands to network and content elements. When the activation is carried out in the network, the billing system should also be notified. Deactivation, modification, and cancellation are usually supported in the same manner.

When activating consumer services at content and/or applications providers, an activation system will handle the command sessions with the servers running applications such as authentication, authorization, and accounting (AAA); Web; mail; domain name server (DNS); music; and games. When provisioning mobile Internet access, an activation system could be used as a mediation layer to communicate with element managers or directly with the network devices.

A single service-activation system that interfaces with the network elements in multiple network technologies will simplify the implementation of support for bundled services. For instance, a service package including 500 minutes of voice service and GPRS access, plus one or multiple content services, could then be handled as a single activation request. Because new services are continuously introduced in the network, there is a strong requirement for activation systems that are easily adapted to support new service logic and new versions of network elements without operational disruptions.

The Critical Need for Business Intelligence

Not withstanding the challenges facing wireless carriers as they introduce mobile Internet access, it becomes obvious that there is a need to carefully monitor business metrics such as customer usage, service profitability, and capacity utilization. During periods of heavy investment in new network technologies and spectrum licenses—compounded by extreme pressures to be profitable—early identification of profitable services now commands the attention of a much wider business audience.

As wireless carriers and service providers introduce new IP–based offerings, they will need to examine and analyze the uptake of such offerings. As addressed previously, carriers and content providers cannot predict which services and applications users will find valuable. Supply-chain players are going to have to experiment with several promotional offerings and closely monitor usage behavior. With real-time data on customer behavior and to what degree promotional services are performing, service providers will have the means to become *proactive* rather than *reactive* with their customers, which can minimize churn and maximize market-share. And that means increased revenue potential.

A business intelligence reporting tool must be easily adapted to a carrier's specific business requirements—current as well as future needs. The business intelligence application should, at the very least, monitor and simplify the analysis of network performance.

High-speed wireless data access technologies such as GPRS and 1xRTT enable "always-on" connection to the packet-switched network, unlike cellular digital packet data (CDPD) or first-generation wireless application protocol (WAP) phones, where users must initiate a connection to the IP network. The notion of "always on" must be explored to its fullest, as location-based services will potentially add significant revenue streams to wireless carriers. However, this opportunity will come at some cost and risk. Although the technology allows for "always-on" connection, actual coverage and the number of cell sites combined with roaming limitations will impede on this concept. Hence, it becomes particularly important to monitor the quality of service (QoS) delivered and relate it to the service-level agreement (SLA) promised. With a powerful business intelligence module, carriers can set SLA violation triggers that can automatically notify the rating engine or billing system to calculate discounts based on the pricing policy contained within the SLA. Furthermore, such applications can effectively assist in forecasting service usage as well as network or capacity expansion planning.

Share Revenue with New Partners

Advertising has been the primary sponsor of content in the fixed Internet world. Generic content has, to a large degree, been free to visitors because advertisers have subsidized the cost of producing the content in exchange for banner ads.

Advertising models will certainly play a positive role in the wireless Internet space; however, they have to be carefully evaluated and positioned for the new medium, namely small handset devices with a very limited display size and much lower resolution than a modern desktop or laptop monitor. Advertising must also be linked to the audience in a much more sophisticated manner. It needs to be connected to the end users via content, locations, and/or m-commerce. In the long term, advertising can still prove to be a viable revenue source for carriers, but until the arrival of full-fledged location-based push and pull services, carriers should focus on offering premium, value-added content and applications that are appealing to customers.

Typically, mobile operators' core business focus is running and managing voice traffic over wireless networks. With the advent of Internet access via mobile phones, an entirely new world of business opportunities emerges. For example, a mobile operator providing access to, say, Amazon.com, for example, is essentially also acting as a point-of-purchase agent. The carrier enables Amazon.com to receive orders from mobile customers that may not have taken place otherwise. In this scenario, the mobile operator becomes a valuable player in Amazon.com's value chain and potentially entitled to be compensated for generating new revenue. By inferring this philosophy to the banking and financial world, the revenue opportunities become enormous—assuming that privacy and security concerns are resolved.

It is crucial that wireless carriers are able to monitor and track the transactions and events generated within their networks and to automate the process of settlements—and revenue sharing—with on-line merchants and content providers. Voice traffic and airtime will most likely continue to account for the bulk of carrier revenues in the short term. However, there are already many different business models emerging in the marketplace, and carriers must start implementing new marketing and service-offering strategies.

The traditional rating function has been to apply a tariff or a retail price to a voice call, where metrics such as duration, time and day of the week, and distance have determined the final charge for the service. As deregulation opened up the market for competition and the increase in capacity that digital networks provided, the wireless industry witnessed the introduction of many creative offerings aimed at targeting various segments with different volume discounts. These "buckets of minutes" offerings required some changes to the billing systems, but these were fairly simple to implement because the main functionality remained the same, namely keeping count of the minutes spent on the network.

With the introduction of IP–based services, carriers must find ways to maximize the revenue potential. To do so, carriers are advised to involve new players and strategic partners in their value chains to develop and create rich content and application services that can be delivered over their wireless infrastructures. As many industry analysts predict, the wireless industry of tomorrow will be characterized by cross-industry partnerships and collaborations, massive reorganization of the value chain, and a competitive environment that rewards innovation and differentiation

The fundamental challenge essentially is how to treat the events occurring in IP networks. Next-generation services, such as high-quality videoconferencing, CD–quality audio streaming, and wireless productivity applications, will

FIGURE 3

The Billing Process

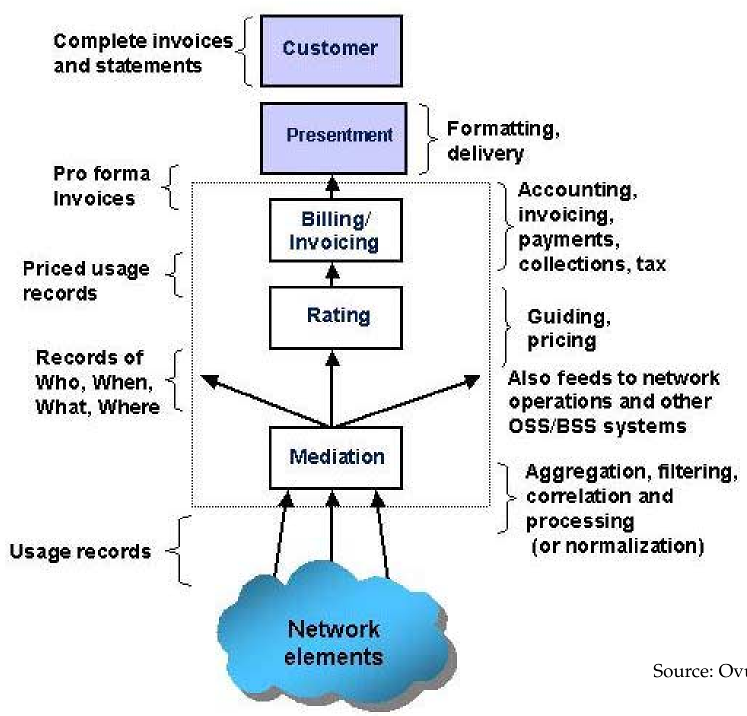

Source: Ovum

require more bandwidth and network resources than, for example, an e-mail with a small attachment. Therefore, it becomes crucial for service providers to ensure that network resource utilization or consumption is compensated for or recovered by the actual users of these bandwidth-hungry applications. As these types of premium services evolve and become part of the offering, new pricing strategies also must evolve. The new billing models must reflect the perceived value of these types of services in order to maximize the revenue potential and achieve sustainable, long-term profitability.

The following are a few emerging revenue-sharing scenarios that the rating component must accommodate:

- Shared revenue
- Sponsored revenue
- Referral revenue
- Billing on behalf of others

A next-generation rating application must be very agile and flexible to support ongoing changes expected in the near future. As already implied, carriers have limited insights of the potential demand and perceived value of these prospective services offerings and therefore will require systems that can easily change business logics, pricing and discount plans, service bundles, limited promotional offerings, and so on. The rating module must also feature the ability to rate events or service delivery on any perceivable dimensions, such as per use, data volume, QoS, number of clicks, and more. To better illustrate the complexity, consider the scenario in the following paragraph.

A business traveler is staying in a city that he/she has not yet visited. In the morning, the individual decides that a cup of Starbuck's coffee is a "must" to start the day. The phone book is full of Starbuck's locations, but the street addresses are unfamiliar. The concierge desk is not yet open, and the staff at the front desk is busy checking out long lines of guests anxious to depart for the airport. However, using his/her mobile phone, finding the closest Starbuck's and receiving directions guiding him/her to that craved cup of coffee is quick, easy, and relatively inexpensive—he/she is charged a nominal 50 cents for this service. Also, consider how the revenue streams can be divided by the parties involved in this information request. For example, Starbuck's may offer the driving directions free of charge to the requestor, while the requestor pays for only the incoming data. Alternatively, Starbucks might sponsor 25 percent or even 50 percent of the information request, including the cost of delivering the data; or they could pay for all of it if the requestor actually went to the coffee shop and spent more than a specified amount of money, say $3. In addition, any of these may have a limited promotional time offer attached to it. This is exciting news for the business traveler and the mobile operator—adding value to voice services by providing content-related Internet services. The business traveler receives a value-added service, and the mobile operator opens a new revenue stream.

Such a simple scenario delivers real value to the end user who doesn't really care how the revenue streams are divided as long as the end user receives the information that he or she is requesting. However, it is a very complicated transaction from a carrier's point of view. Managing such relationships and transactions is a brand new domain for mobile operators, and their current support systems are not intended for such functionality.

As the scenario illustrates, a rating engine needs to be extremely flexible and adaptive to accommodate such fast-changing business logic. The engine must enable carriers to price events and transactions on multiple dimensions as well as the ability to split and calculate the revenue streams—both at the retail and wholesale levels. By having well-defined interfaces based on industry standards such common object request broker architecture (CORBA), extensible markup language (XML), and Java, carriers have the opportunity to cost-effectively integrate a powerful standalone rating component to their existing billing

FIGURE 4

Split Rating for a Single Traffic Case to Allow for Revenue Sharing

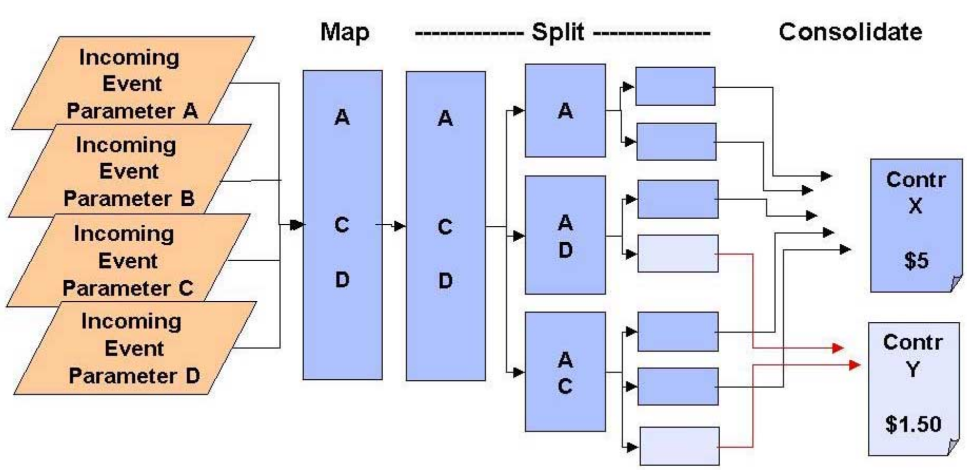

system. Even if customized interfaces are required, such an approach could yield a favorable ROI scenario compared to re-engineering efforts or replacement of the existing billing system.

The main objective with rating module architecture is to effectively handle the rating of any conceivable service without the need for additional programming. All rating characteristics should be managed through set-up definitions of the business rules that a service provider wants to apply, without any customization needed. Furthermore, the rating module must be capable handling extreme performance requirements. Even with the advanced analysis, consolidation, and processing functions, hundreds of millions of daily event records may reach the rating engine, which should be able to handle such a massive volume in a single-server environment.

Conclusion

To gain and sustain a competitive edge, wireless carriers must take a proactive approach to all aspects of their operations. Their long-term success is in the marketplace is largely dependent on their ability to deliver quality services at prices, which the customers are willing to pay. With the enormous amount of capital being invested in next-generation infrastructure and spectrum rights, carriers are simply forced to roll out new services and products in record speed to an increasingly sophisticated and demanding customer base. These new services will not be delivered in the traditional supply chain to which mobile operators have been accustomed but rather through a dynamic value chain consisting of a multitude of players, which dictates entirely new business models and partnership strategies for carriers to implement and execute.

In an industry where change is the only constant, Ericsson Customer-Management Solutions (CMS) for mobile Internet can help to make the difference in ensuring accurate revenue accounting and sharing between multiple partners and seizing the new opportunities that m-commerce will provide. Ericsson CMS can help wireless service providers to capitalize on these opportunities through extensive knowledge and solutions that protect existing system investments; accelerate new market launches; and capture valuable customer, service revenue, and network intelligence while enabling revenue sharing with new partners.

Notes

1. Adam Zawel, "Wireless Internet: How Much Money Will Consumers Really Pay?" The Yankee Group Report, Wireless/Mobile Services Report, Vol. 1, No. 17 (November 2000).
2. Merrill Lynch – Global Securities Research & Economics Group, "The Next Generation V," Wireless in the U.S. (March 9, 2001).

Deploying Voice Portals in a Network

Michael Krasner

Vice President and General Manager, Voice Solutions

Comverse

1. An Introduction to Voice Portals

The word *portal* comes from Medieval Latin and means "city gate." Through a portal comes access to all that is within. Yahoo!® used the word portal to describe its one-stop directory service to nearly everything the Web has to offer. Voice portals are a logical extension of Web portals. Simply defined, a *voice portal* allows users access to information, check voice, e-mail messages, and place calls over the telephone using nothing but their voices. The objective of a voice portal is to simplify the way people use telephones to communicate, access information, and conduct business.

Voice portals use powerful computers and software to recognize speech commands, find information, listen to audio streams, convert text into an audible format, and help users to keep in touch. A voice portal may be best thought of as a voice version of a Web browser that accepts spoken commands and returns audio data and responses. The potential breadth of commands, applications, and information that can be accessed with a voice portal is limitless. For example, a voice portal can add custom applications, such as a voice equivalent of a contact management program. In addition, voice portals can be integrated into existing telephone networks and aid in the basic user functions of that phone system, such as placing calls and handling voice-mail. Today, voice portals can read e-mail, manage your voice-mail, provide directions, read the latest news and sports scores, play a game, manage your schedule, and much more. In the future, voice portals will book your flights, arrange flower delivery to your loved ones, and serve as your personal concierge.

2. Voice-Portal Drivers

The phenomenal success of the Internet demonstrates our insatiable desire for information. However, accessing the Internet is generally a task that requires sitting down at a computer or some other Internet appliance to fully take advantage of its services. In an attempt to mobilize the Internet experience, mobile phones have started to offer Internet content. Adequately displaying the Internet's data-intensive information has been difficult to accomplish on mobile phones that are constantly shrinking in size. This is further complicated by the increasing demand for safe and hands-free solutions. The only truly mobile and hands-free Internet solution is voice interaction.

The voice marketplace has been driven by several factors that have evolved over the last few years. First has been the development of inexpensive and powerful computer technology, which has led to the development of reliable and reasonably priced voice-recognition systems. Second, speech-recognition applications have matured to recognize a broad array of speech patterns and words, enabling systems to be very user-friendly. Third, the availability of a broad array of information and data that can be transmitted wirelessly, such as e-mail and news feeds, has helped introduce new wireless services. Fourth, the end-user cost of entry is very low because voice-portal technology is transmitted completely via voice, which means any phone can access a voice portal, regardless of the phone's brand or technology within the phone. Finally, the amount of time that people spend using mobile phones in their cars has greatly increased, making a hands-free solution a critical safety concern.

3. Voice-Portal Market and Opportunity

According to the technology research firm Ovum, there are nearly 100 million cell phones in use in the United States alone. In 1999, 2.1 billion wireless calls were placed to 411 for information (Strategis Group, 2000). North American voice-portal use among wireless users is expected to mushroom from 1 million users in 2001 to 56 million by the end of 2005 (Allied Business Intelligence, 2000). Another report by the Kelsey Group estimates a conservative 45 million voice-portal users by 2005 in North America and estimates that 18 million users will use voice portals to shop for products (see *Figure 1).*

Estimates for potential revenue vary considerably among analysts because of the many possible models for income. The Cahners In-Stat Group predicts at least $1.6 billion in revenue by 2005 through local advertising, national advertising, sponsorships, hosted content, and revenue from third-party sales. They note that significantly higher revenue could be generated—upward of $32 billion—if customer profiles are sold. More aggressive estimates come from the Kelsey Group, which predicts $5.4 billion by 2005 from advertising and commerce revenue only.

4. What to Look for When Deploying a Voice Portal

Deploying a new technology, such as a voice portal in a wireless network, is an elaborate process that requires extensive

FIGURE 1

Number of Voice Portal Shoppers

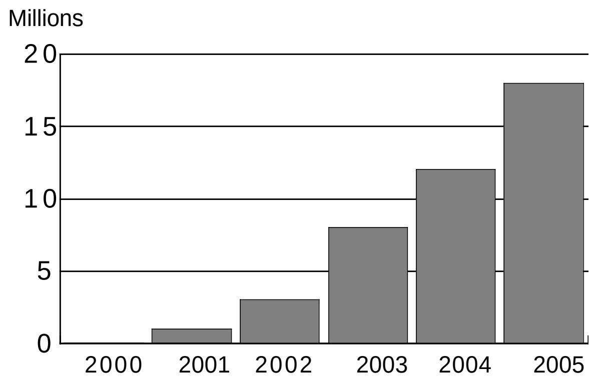

Source: The Kelsey Group, 2000

vendor analysis. The following topics are important issues to consider when looking for a voice-portal vendor.

4.1. Applications

Voice-portal technology is only as good as the breadth and depth of the applications designed to support it. Minimum requirements must include voice-activated dialing and voice-controlled messaging. Voice dialing must accommodate dialing by phone number and by name. Phone number dialing must accommodate both single number speaking (e.g., "one, eight, zero, zero") and shortcut dialing (e.g., "one, eight hundred"). Ideally, voice dialing should be integrated with a personal address book that provides a great deal of flexibility. Some of the more advanced address-book options include a preprogrammed list of major vendors (e.g., major airlines) and allow the user to synchronize his or her personal information manager (e.g., Palm™ Pilot) with the voice portal using a traditional Web interface. Address-book dialing must allow voice dialing with both name and location options. For example, "call Mike at the office" must be distinguished from Mike's home and mobile phone numbers. Additionally, voice-controlled messaging must provide an intuitive and simple interface that enables users to listen, filter, store, send, and delete messages verbally.

More advanced applications include a voice browser for accessing topical content, such as weather, traffic information, and news. A good voice browser must allow content to be customized and localized to each user. For example, a user can indicate the city in which he or she lives and perhaps news topics of personal interest, providing customized content not available over alternative communication mediums, such as a radio. Some more advanced browsers allow users to preset their interests using a computer and a Web browser. In addition, live content is critical to the user experience. With a voice browser, the phone becomes a mobile Web radio and enhances traditional news radio content.

4.2. Language and Human Factors

For global telecommunications companies, it is important to implement new technology localized to different cultures and that supports different languages and dialects. When implementing voice-portal technology, select a vendor that can deploy internationally and has experience with system implementation and language localization. In addition, users should be able to profile their language of choice in environments where multilingual support is critical.

In addition to language considerations, a user interface that is natural and creates a positive experience—also known as human factors—is critical to wide-scale consumer adoption of a voice portal. A good voice-portal vendor must understand the limitations of the mobile telephone channel and design voice commands that do not confuse the portal yet are very user-friendly. The science of human factors is not well-understood by most voice-portal vendors and is critical to the overall user experience. Similar to a conversation with a live person, a voice portal must know how to handle idiosyncrasies common in speech interactions. For example, the ability for a user to interrupt the voice portal at any time and initiate new commands is critical to the user experience. A good voice portal will know how to smoothly handle such interruptions.

Verbal commands that are easy to remember and speak are also critical to the user experience. A more basic voice portal will constantly review a list of options and sub-options that users have, similar to a voice-mail system. The problem with this approach is that it is complicated, time-consuming and not the way people speak in natural situations. A natural-language interface must use short prompts that are easily controlled by the user.

4.3. Branding

An essential factor to consider when implementing a voice portal is how easy it is to develop a recognizable brand for the service. This user experience is accomplished primarily via the personality of the voice portal. Much like a person, a voice portal should have a unique demeanor communicated by its voice, tempo, and language formality. The personality style should be conveyed by prompts and the grammar of the voice portal. Possible options may include the following:

- ***Impersonal:*** "Mailbox action?"
- ***Trendy:*** "I got your mailbox open. What now?"
- ***Friendly and Informal:*** "What should I do with your mailbox?"
- ***Professional:*** "I have opened your mailbox. How shall I proceed?"

Coupling the above grammars with verbal personalities—such as a robotic voice, a young person, a happy person, or a butler—will reinforce the brand experience. Personality creation is a complex yet important process that requires a thorough analysis of your target audience and an understanding of what type of verbiage will elicit expected results with users.

Advanced voice-portal vendors will allow customized demographic content for your various audiences. For example, the list of preprogrammed vendors in a public voice directory may be tailored for specific age groups. A teenager profile might include local pizza, arcade, and movie-theater listings. Additional branding considerations should include customized help message content, modifiable prompts, unique chimes, and grammar modifications.

4.4. Customization and Scalability
A voice portal should integrate with your existing mobile telephone network without requiring you to adopt proprietary technology. An advanced voice portal will be customizable to comply with the varying standards and billing systems used by the telephone industry. With the incredible growth of mobile users, the voice-portal platform must also scale to accommodate up to millions of users and support distributed architecture (e.g., include support for roaming). It is also important to integrate third-party directory and Web servers via standards such as lightweight directory access protocol (LDAP) and VoiceXML.

4.5. Speech Recognition and Text-to-Speech
The ideal voice portal is not limited to a single vendor for speech-recognition or text-to-speech capabilities. With the evolving software techniques for recognizing speech and converting it to text, a voice-portal vendor should have a flexible system that allows new technology from other vendors to be integrated with the system in the future.

4.6. Deployment
When deploying a voice portal, it is important to work with a vendor that has established a systematic procedure for the service launch. Issues to consider include the following:

- ***Internal Testing:*** This phase should test general functionality with a small sample of users and identify areas that need to be optimized or corrected.
- ***External Testing:*** This type of testing should be conducted on a large scale and includes testing acoustic models, voice-recognition parameters, grammar quality, and ease of use.
- ***Off-Line Tweaking and Internal Lab Testing:*** This phase should include testing the voice portal's listening capabilities and assessing if the voice portal became confused or erred in any transactions.
- ***Soft Launch:*** During this phase, the voice portal should be integrated with existing voice-mail systems, and support personnel should be trained on the hardware, platform, administration tools and the nature of speech systems. In addition, a Web interface may be aligned with the voice portal at this stage.
- ***Hard Launch:*** This phase is the official launch, where marketing collateral would be released and the entire system would be available for use. An ideal voice-portal vendor should be able to aid in every one of these phases of deployment.
- ***Follow-Up:*** This recommended step will allow for future refinement of the system and help gauge the user experience.

4.7. Vendor Experience
When selecting a voice-portal vendor, it is important to assess their track record and experience relative to your marketplace. A voice-portal vendor should be well established and stand by its claims.

5. Summary

From the beginning of time, no communication method has been more important than the spoken word. It is no surprise that we seek to enhance our verbal communication with everything from answering machines to mobile phones. Modern technology has removed the physical boundaries of communication and helped to create a mobile society.

The next step for voice technology is to extend the flexibility and functionality of the Internet—fundamentally a visual medium—to voice-driven applications. Today, this can best be accomplished by voice portals—speech-driven applications that listen, respond, read, and record commands and content. Voice-portal technology is best applied to contexts in which normal Internet access is restricted or nonexistent. The largest business opportunity lies in the mobile marketplace. With the wide adoption of cellular phones, mobile consumers are the ideal target audience for speech portals.

MIMO: The Preferred Choice for Last-Mile Access

Rajeev Krishnamoorthy
Vice President, Technology Development
Iospan Wireless Inc.

This article addresses some of the fundamental technical challenges related to mass deployment of next-generation fixed wireless access (FWA) systems. It presents and discusses state-of-the-art wireless multiple-in, multiple-out (MIMO) technology recently developed to overcome the perennial challenges of line-of-sight dependency, spectral inefficiency, and multipath and installation difficulties.

1. Technical Challenges for FWA

The demand for broadband Internet access is booming. Delays in the deployment of 3G high-speed wireless networks (wideband code division multiple access [WCDMA] and CDMA2000 are unlikely to be widely available before 2003 or 2004) and slow progress in satisfying demand for wireline access place extraordinary expectations on alternative last-mile technologies such as FWA. This is especially apparent for the lower bands (multichannel multipoint distribution system [MMDS] or 2.6 GHz, and 3.5GHz) dedicated to mass-market residential, small office/home office (SOHO) use. *Figure 1* illustrates the main network components of FWA.

Despite strong customer demand, current rollout of first-generation FWA equipment by service providers such as Sprint and WorldCom is stumbling due to fundamental technical limitations, including line-of-sight dependency and costly, time-consuming subscriber antenna installation. In addition, it is widely believed that current FWA technology cannot meet the forecasted demand for capacity because of its inefficient way of using of spectrum. To successfully respond to demand and to carry their future share of the access traffic, it is necessary for the second generation of FWA systems to deal with the following requirements:

1.1. High Performance

The demand for high capacity and the competition with landline (digital subscriber line [DSL] and cable modems) and upcoming wireless data services (2.5G and 3G) make high performance a premium requirement for FWA networks. This is especially true, given that these networks will not support mobility. Superior performance of FWA over competing access solutions can be achieved in three key areas:

- Aggregate user rates
- Spectrum efficiency
- Coverage

High aggregate data rates must be offered in order to support rich multimedia services, in both downlink and uplink. The aggregate rate is the rate that a single user will experience averaged over all possible subscriber locations.

Because spectrum is a scarce resource, especially in the sub-5 GHz band, it must be used efficiently. The spectrum efficiency can be defined as that ratio of aggregate rate to the product of bandwidth used and frequency reuse factor. To be efficient, 2G FWA must reuse frequency aggressively, and, as such, the transmission technology should be made tolerant to interference.

To get the most of every installed base transceiver station (BTS) in initial rollouts, the coverage performance is critical. This can be measured by the cell radius served by the base, with a constraint on the ratio of home locations that must be covered, typically 90 percent.

Finally, link latency and reliability are also important measures, and the performance there should be comparable to existing wired access technologies. Overall, the typical performance requirements for 2G FWA systems are shown in *Table 1*.

1.1.1. Unfriendly Channel Environment

Because it is desirable to serve most of the cell area, the dependency on a line of sight at the subscriber's home must be eliminated. In many locations, this results in deep channel fading caused by multipath propagation, similar to what is observed in standard cellular networks. In addition, broadband transmission over multipath channels introduces frequency-selective attenuation, also due to intersymbol interference. Finally, reuse of the frequency in neighboring cells causes co-channel interference at the receiver. In practice, the severity of fading and interference will vary from location to location and time to time. However, due to the stringent link reliability requirement imposed by the service providers, the transmission technology must cope with the worst-case channels at all locations

TABLE 1

Performance Requirements for 2G BWA

Figure of Merit	Requirement
Aggregate Rates	4 Mbps
Spectrum Efficiency	1 Bit/Hz/BTS
Coverage	3 Miles (90% area)
Latency	Comparable to DSL
Link Reliability	0.999

1.2. High Level of User Friendliness

In addition to high performance requirements, the success of FWA rollouts depends on the ability of the service provider to quickly install the equipment. Current technology requires skilled, accurate pointing of the rooftop antenna toward the line-of-sight direction. This results both in many unsuccessful installation trials (since it is difficult to predict the existence or not of a line of sight at the subscriber location) and costly, difficult installations at those locations that have line of sight. Ideally, of course, the subscriber's unit is self-installable, with no constraint on the room/location where the antenna will be put. As a nearer goal, however, 2G FWA networks can be made much more flexible than current systems by the use of low-height, compact-size, wide-beam subscriber antennas.

Unfortunately, the use of user-friendly antennas as described often means further degradation in the channel characteristics, because this type of installation introduces higher path loss and more multipath, as well as increased interference. The ability for the modem to operate at low levels of signal to noise (SNR) and interference ratios (SINR) is therefore essential.

1.3. Low Deployment and Maintenance Cost

For the subscriber unit, the cost should be comparable to xDSL and cable modem (while presumably offering higher performance). On the network side, in addition to the fast deployment that is inherent to wireless, significant cost savings can be gained by having a high coverage performance in initial rollouts and high spectrum efficiency in mature deployments with high subscribership. In all cases, the cost of deployment is tightly coupled to the network performance.

2. Design Features for 2G FWA

To meet the challenge of high data rates, high coverage, and high spectrum efficiency in an unfriendly propagation environment, state-of-the-art wireless research and techniques can be applied. In this article, we touch on selected technologies addressing this problem. We distinguish the physical layer (Layer 1) from the link layer (Layer 2), although a successful design calls for layer integration.

For Layer 1, key design features include antenna diversity, MIMO processing and spatial multiplexing, adaptive interference canceling, orthogonal frequency division multiplexing (OFDM), and intelligent rate adaptation mechanisms. For Layer 2, we focus on the link-level retransmission mechanism and fragmentation (ARQ/F).

2.1. Physical Layer

2.1.1 Spatial Diversity

Spatial diversity (SD) is obtained through the use of multi-element antennas at either the BTS and/or the customer-premises equipment (CPE). Antenna combining can be used to deal efficiently with multipath fading. The basic idea of antenna or SD is that several uncorrelated fading signal are much less likely to fade simultaneously than a single one, hence intelligent combining of the antennas can offer a much more stable link for the upper layers to work with (see

FIGURE 1

FWA Network Components

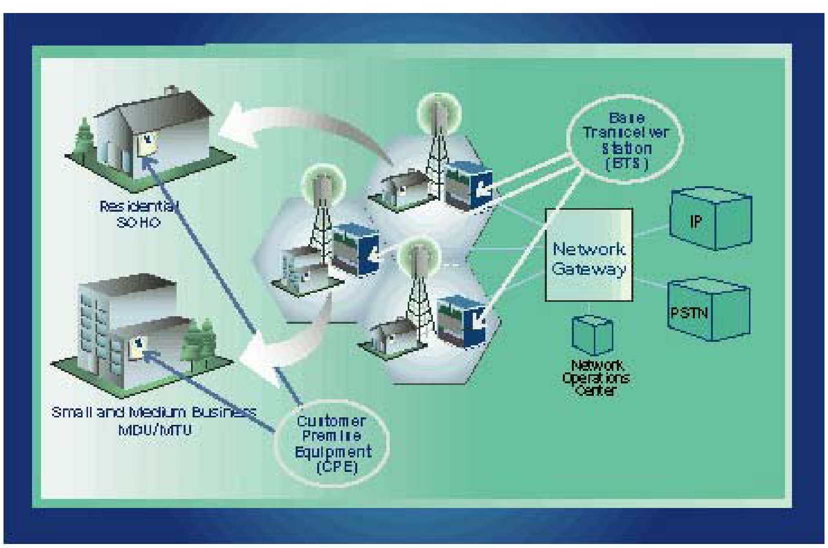

Figure 2). In addition, SD is a robust form of diversity, because low antenna correlation (less than 0.7) is almost always guaranteed with sufficiency antenna spacing. Typically, one wavelength of separation will suffice at the subscriber's side, for instance. Without diversity, the transmitter is obliged to operate with large margins of power for fading to achieve a target packet-error rate (PER), thus resulting in a much smaller coverage compared to a system with SD. In stationary wireless Internet protocol (IP) access, the ability to combat fading is even more critical, due the fact that fades last typically longer than in mobile links (up to a few seconds), which can cause the transmission control protocol (TCP)/IP connection to reset its congestion window and result in unacceptable delays.

Use of SD can reduce the SNR requirement by 10 to 15 dB with no loss of TCP/IP performance. This can be exploited not only to extend the coverage, but also to increase data rates and to allow a tight frequency reuse that will work in arbitrary terrain (flat or hilly). Antenna decorrelation can also be realized by using orthogonal polarization states of the antenna elements.

2.1.2 MIMO Diversity and Spatial Multiplexing

FWA can benefit much by using multi-element antennas at both ends of the link (so-called MIMO). MIMO can be used for enhanced SD because transmit *and* receive antenna elements can be jointly combined to combat fading. However, the most spectacular benefit of MIMO is a form of transmission called spatial multiplexing (SM). In MIMO–SM, the data rate is literally multiplied, at no cost in bandwidth or power, by exploiting simultaneous transmission over the so-called eigen-modes of the (matrix) channel [Paulraj94, Foschini96]. In practice, this means that several transmit antennas can send independent signals over the same frequency and time slots. The bits are mixed through the channel but are recovered through spatial and temporal processing at the array of receive antennas, as shown in *Figure 3*. The gain in data rate with SM is proportional to the minimum of the number of transmit and receive antenna. A key backer of this technology, Iospan, has developed specific algorithms to optimize this technique under various channel conditions. MIMO–SM is particularly well suited for FWA because, unlike handsets, compact multi-element subscriber units can easily be designed with two or three elements. For example, for a two-transmitter two-receiver system, the peak data rate is doubled.

2.1.3 Adaptive Interference Canceling

Adaptive interference canceling (AIC) is another form of spatial combining technique that aims at mitigating interference originating from neighboring cells for the purpose of reusing the frequency more efficiently. AIC exploits the idea that the desired signal and the interfering signal do not usually have the same spatial signature. Therefore, by placing a "null" in the direction of the interferer's signature, it is possible to suppress it.

2.1.4 Frequency Diversity in OFDM

In traditional wireless systems, the problem of inter-symbol interference tackled through the use of equalizers—i.e., time-domain filters implemented at the receiver that seek to invert the channel multipath response. Although easier to implement, one disadvantage of this approach is that the complexity of such filters grows exponentially or at best linearly, with the bandwidth making it unsuitable for future very-high–rate systems. Recently, competing multicarrier approaches, such as OFDM modulation, have gained popularity because of their robustness. OFDM works by splitting the original high-rate bit stream into low-rate substreams,

FIGURE 2

Multiple Antennas Provide a Robust Source of Diversity in Non–Line-of-Sight Channels

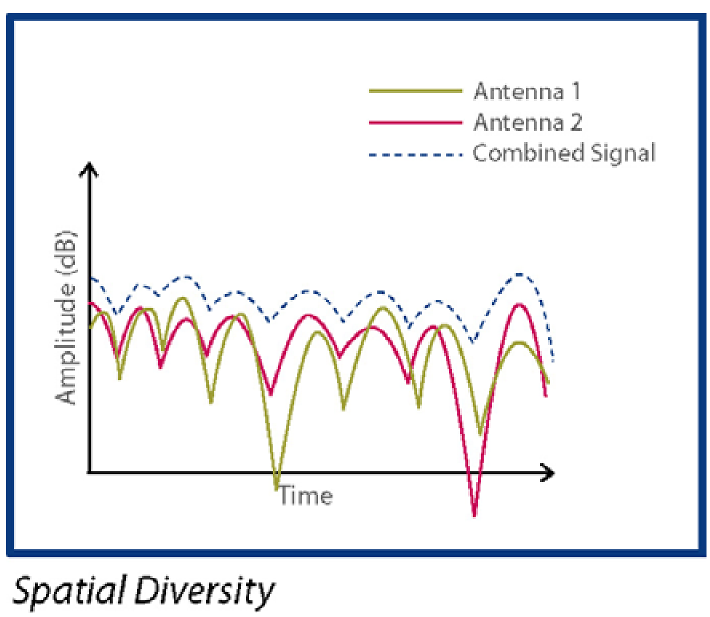

FIGURE 3

Simplified Spatial Multiplexing Scheme with Three Antennas
This gives three-fold data-rate improvement.

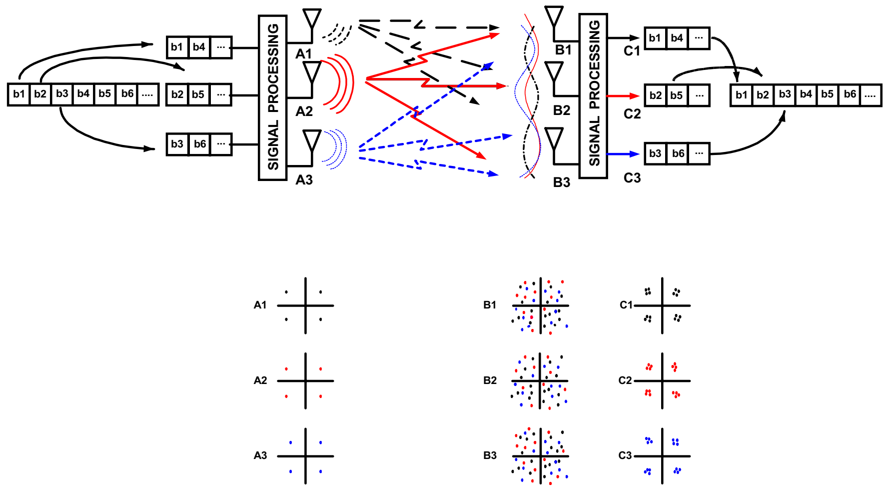

each mapped onto one of many adjacent subcarriers. Each subcarrier is small enough to remove the need for equalization. In practice, the OFDM is implemented efficiently using fast Fourier transforms. One bonus of the OFDM approach is that is effectively converts multipath into a benefit through the idea of frequency diversity, as shown in *Figure 4*. This can be realized by overlaying coding techniques on top of the OFDM modulation, which has the effect of spreading information bits equally over the entire modulation band. If the delay spread is large and the fading is frequency-selective, the subcarriers with good SNR help to mitigate the effect of fading occurring at other subcarriers, making the channel effectively look like a non (or less) fading one. Such techniques lead to several decibels (dB) of gain in the link budget, extending coverage and improving reuse performance.

2.1.5 Adaptive Modulation

Adaptive modulation (AM) and coding is a key technique for enhancing the spectrum efficiency of FWA. AM has been first popularized in the enhanced data rates for GSM evolution (EDGE) cellular standard. AM lets a user adapt its data rate as a function of channel conditions (e.g., SINR, fading rate). In doing so, it gives two important benefits. First, it makes CPE installation easier and more forgiving, because mispointing of the antennas can be compensated automatically by switching to a lower modulation and stronger code at the physical layer. Also. Edge-of-cell users can be more easily covered by these low-rate modes of transmission. Second, the overall spectrum efficiency is improved several-fold because all subscribers who enjoy better than edge-of-cell channel conditions can use their margin of SNR to increase their rate proportionally. By comparison, a non-adaptive system will have to be deployed with the most conservative modulation/coding for all users in order to preserve good coverage and frequency reuse.

Furthermore, a reasonably fast adaptation rate (every 100 ms or so) will allow the modem to track the time variation, such as fast fading of the channel at the subscriber's location. This feature is especially important to deal with unforeseen change of line-of-sight conditions that can occur due certain obstacles (trucks, tree leaves, etc.). In the case of OFDM modulation, advanced algorithms that track channel characteristics in the time and frequency domain can be developed.

The adaptation algorithm can also be designed to react to spatial properties of the channel so as to optimize not only the modulation/coding level, but also the way in which input signals are mapped onto the antennas, again for the sake of maximizing data rates. For example, if the link is in a SM mode and the number of usable eigen-modes suddenly collapses due to adverse propagation conditions, the system falls back to a MIMO diversity scheme, and vice versa.

2.2. MAC and Link Adaptation Layer

2.2.1 Automatic Retransmission Request/Fragmentation (ARQ/F)

Despite the use of coding and diversity techniques at the physical layer, it is virtually impossible to hide all channel impairments from the protocol layers. In fact, a sensible system design of FWA calls for a physical layer that operates at a high level of PER (more than 1 percent) to minimize the SINR requirements at the receiver, making it possible to extend the coverage and tolerate aggressive frequency reuse. To achieve high TCP/IP throughput, on the other hand, excellent IP PER (less than 0.1 percent) must be passed on to the network layer by the link layer. This can be done using automatic retransmission request/fragmentation (ARQ/F) algorithms. ARQ/F consists of implementing an acknowledgment and retransmission mechanism at the link layer between the subscriber unit and the BTS. The media access control (MAC) layer performs IP fragmentation over the air into "atomic" data units (ADU), so only errored ADUs are retransmitted. A technique complementary to channel coding, ARQ/F has the merit of introducing redundancy only during the small fraction of time when data gets corrupted. A finite bufferization mechanism makes ARQ/F efficient in dealing with short as well as long

FIGURE 4

OFDM with Coding Transforms Multipath into a Source of Anti-Fading Diversity

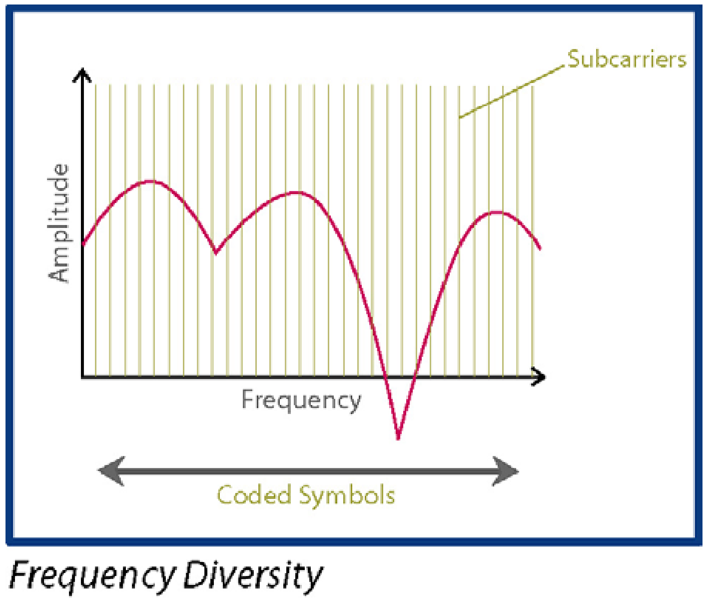

fades. ARQ/F removes packet errors at the cost of only moderate additional link latency. For this reason, ARQ/F–based systems can be designed to operate at high PER levels while still providing satisfactory quality of service (QoS). The higher operating PER will routinely translate into SINR gains of up to 6 dB in the link budget.

3. Conclusion

Successful deployment of 2G FWA networks poses important technical but solvable challenges. Chief among those is the requirement to perform at a high level of performance in an unfriendly wireless channel environment. Multipath mitigation techniques such as antenna diversity and OFDM modulation appear as essential design elements of FWA. Enhanced processing techniques such as MIMO–SM and adaptative modulation can give the desired higher data-rate performance. Finally, interference canceling and retransmission mechanisms will provide enhanced robustness for higher frequency reuse and coverage.

Implementing Distributed Look-Up Services in Wireless Ad-Hoc Sensor Networks

Alvin Lim
Professor, Department of Computer Science and Engineering
Auburn University

Abstract[1]

Wireless sensor networks are becoming increasingly important in the implementation of pervasive computing systems. Real-time information can be gathered directly through a very large number of mobile sensor devices scattered over a wide area and disseminated to various applications and end users over the ad-hoc sensor network with little or no fixed network support. Distributed computation, such as collaborative target tracking and remote surveillance, can be done on the fly as the sensor data are disseminated through the sensor network. These large surveillance sensor networks must adapt rapidly to dynamic changes in sensor node configuration. Appropriate distributed services and network protocols are required to solve the problems of mobility, dynamic reconfiguration, dispersion, weak and intermittent disconnection, and limited power availability. We provide three main distributed services to deal with these issues: look-up service, composition service, and dynamic adaptation service. In this paper, we focus on the implementation of the distributed look-up service that manages changes in the available services in the sensor network, migration of sensor nodes, and changes in task and network requirements. The distributed look-up service is implemented using directed diffusion network. The distributed look-up service provides more efficient and large-scale support for these changes than would be possible using directed diffusion alone. Through a distributed implementation of the look-up services, other application-specific network and system services can be defined spontaneously in the sensor network.

1. Introduction

Real-time information gathering is critical in many dynamically changing enterprise systems, such as the battlefield and commercial inventory and distribution systems. This real-time information can be efficiently gathered and disseminated via integrated low-powered sensors and mobile devices [11,19,21] that are deployed throughout the enterprise pervasive computing system. These mobile and miniaturized information devices are equipped with embedded processors, wireless communication circuitry, information storage capability, smart sensors, and actuators. These sensor nodes are networked in an ad-hoc way, with little or no fixed network support, to provide the surveillance and targeting information for dynamic control of the enterprise. Sensor devices are mobile, subject to failure, deployed spontaneously, and repositioned for more accurate surveillance. Despite these dynamic changes in the configuration of the sensor network, critical real-time information must still be disseminated dynamically from mobile sensor data sources through the self-organizing network infrastructure to the components that control dynamic re-planning and re-optimization of the theater of operations based on newly available information.

With a large number of sensor devices being quickly and flexibly deployed in most impromptu networks, each sensor device must be autonomous and capable of organizing itself in the overall community of sensors to perform coordinated activities with global objectives. When spontaneously placed together in an environment, these sensor nodes should immediately know about the capabilities and functions of other sensor nodes and work together as a community system to perform cooperative tasks and networking functionalities. Sensor networks need to be self-organizing, because they are often formed spontaneously from a large number of mixed types of nodes and may undergo frequent configuration changes. Some sensor nodes may provide networking and system services and resources to other sensor nodes. Others may detect the presence of these nodes and request services from them.

Other challenges in implementing ad-hoc sensor networks are as follows: First, many different types of sensors with a range of capabilities may be deployed with different specialized network protocols and application requirements. Data-centric network protocols are becoming common in sensor networks [5,6]. With many mixed types of sensors and applications, sensor networks may need to support several data-centric network protocols simultaneously. Second, these mixed types of sensor nodes may be deployed incrementally and spontaneously with little or no pre-planning. Third, the sensor network must react rapidly to changes in the sensors' composition, task or network requirements, device failure and degradation, and mobility of sensor nodes. Finally, distributed services and network protocols for a sensor network must address the problems of weak connection, intermittent disconnection, and limited power availability.

The underlying principle for building self-organizing sensor networks that can overcome the aforementioned challenges is to provide the fundamental mechanisms upon which other networking and system services may be spontaneously specified and reconfigured. The three fundamental mechanisms are service look-up, sensor-node composition, and dynamic adaptation. Through a distributed implementation of these look-up servers, composition servers, and adaptation servers, other network and system services can be defined spontaneously in the sensor network. They also dynamically adapt these services in response to device failure and degradation, movement of sensor nodes, and changes in task and network requirements.

Application-specific network and systems services may be provided impromptu by sensor nodes and supporting nodes, including application-specific information dissemination and aggregation services. A sensor network application may utilize negotiation methods [10] for information dissemination. Sensor nodes will instead register these services with the look-up server.

2. Architecture

We use an approach that integrates three self-configuring system layers: sensor-application systems, configurable distributed services, and the network and physical-device layer. The sensor-application layer contains self-configuring sensor applications, collaborative signal processing, and adaptive sensor query processing. The runtime reconfigurable distributed system contains appropriate distributed services for supporting self-configuring sensor applications. The network and physical layer contains data-centric network routing protocols, physical wireless transmission modules, and sensors that generate the raw data.

The main functions of these key system layers (see *Figure 1*) are as follows:

1. ***Self-organizing application systems***, e.g., sensor information-processing layer and collaborative signal processing
2. ***Configurable distributed systems*** that provide distributed services to the self-configuring application systems
3. ***Sensor-networking and physical-devices layer*** that routes messages through the ad-hoc sensor network and gather real-time sensor data

At the physical-device layer, different physical sensor and mobile devices may be assembled impromptu and reconfigured dynamically in an ad-hoc wireless network. Each sensor node contains a battery power source, wireless communications, multiple sensing modality, a computation unit, and limited memory.

Dual processors may be included for computation and real-time sensor processing. Three common sensing modalities are supported: acoustic sensing using commercial microphones, seismic vibration using geophones, and motion detection using two-pixel infrared imagers. Wireless transceivers in the nodes provide communications between nodes using time division multiplexing (TDM) and frequency hopping spread spectrum. In each cluster, neighboring nodes communicates through a master node that establishes the frequencies used by the nodes. Each node contains a global positioning system (GPS) receiver that allows the node to determine its current location and time.

At the networking layer, ad-hoc routing protocols allow messages to be forwarded through multiple physical clusters of sensor nodes. Directed diffusion routing [6] is used because of its ability to dynamically adapt to changes in sensor network topology and its energy-efficient localized algorithms.

To retrieve sensor information, a node will set up an interest gradient through all the intermediate nodes to the data source. Upon detecting an interest for its data, the source node will transmit its data at the requested rate.

The configurable distributed system uses the diffusion network protocol to route its messages in the presence of dynamic changes in the sensor network. These distributed services support self-configuring sensor applications systems, such as distributed query processing, collaborative signal processing, and other applications. Through the distributed look-up services, application and system programs can be spontaneously and independently fielded and will be able to determine each other's services and self-organize themselves to perform collaborative tasks. Another advantage of the distributed look-up services is that applications programs may use simpler communication interfaces and abstraction than the raw network communication interface and metaphor (e.g., subscribe/publish used in diffusion routing).

Furthermore, these distributed services may enhance the overall performance, such as throughput and delay. These services will be implemented on top of the directed diffusion protocol, which can still be used by applications concurrently with these distributed services. Directed diffusion can still be the preferred method for retrieving sensor data and will be used by some of the services.

At the self-configuring application system layer, distributed sensor query processing and collaborative signal-processing modules communicate with each other to support monitoring, surveillance, and tracking functions of the enterprise. In sensor information systems, the cooperation between mobility-aware mediators, sensor agents, and collaborative signal-processing modules provide efficient access to diverse heterogeneous sensor data, surveillance, and tracking information through the sensor network. The mobile sensor information layer is supported by three major components: interoperable mobile object, dynamic query processing, and mobile transactions. In the interoperable mobile object

FIGURE 1

Architecture of Self-Organizing Sensor Networks

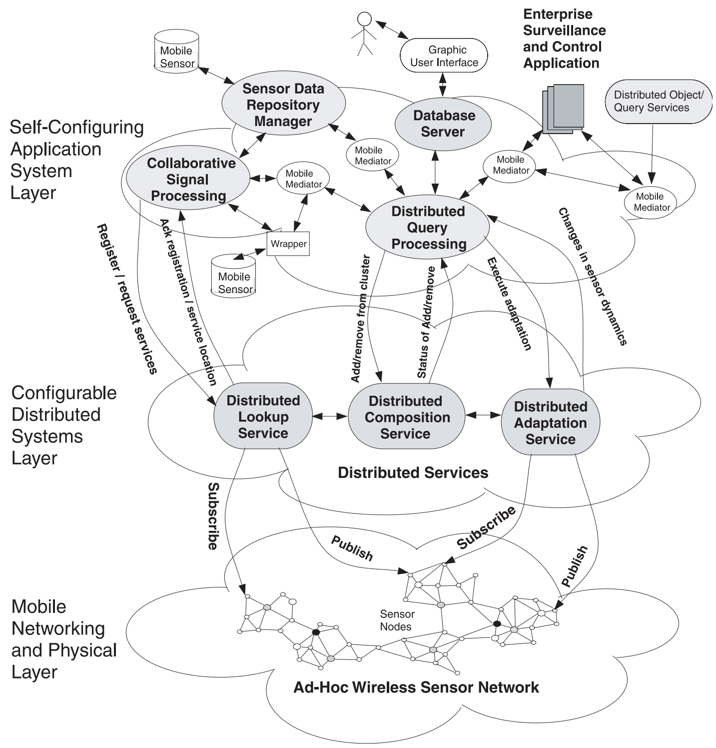

model, a cooperative network of mobility-aware mediators and sensor agents will be configured to support interfaces to remote sensor data sources through multi-hop wireless network protocols.

3. Data-Centric Network Protocols

We use directed diffusion protocol [8] to implement all of the distributed services and for retrieval of data through dynamically changing ad-hoc sensor networks. Diffusion routing converges quickly to network topological changes, conserves mobile sensor energy, and reduces the network bandwidth overhead, since routing information is not periodically advertised. Routing is based on the data contained in sensor nodes rather than unique identification. Directed diffusion is a type of reactive routing protocols that only updates routing information on demand. In contrast, proactive routing protocols, such as links-state routing, frequently exchange routing information. For sensor networks that experience greater dynamic changes, reactive routing algorithms are more appropriate, whereas for those that are more static and experience infrequent topological change, proactive routing algorithms are more efficient.

Directed diffusion is a data-centric protocol, i.e., nodes are not addressed by Internet protocol (IP) addresses but by the data that they generate. Data that are generated by a node are identified by their attribute-value pair. A sink node requests for a certain data by broadcasting an *interest* for the named data in the sensor network. The interest and gradient is established at intermediate nodes for this request throughout the sensor network. When a source node has a data that matches the interest, the data will be "drawn"' down toward that sink node using this interest gradient that was established. Intermediate nodes may cache, transform data, or direct interests based on previously cached data. The sink node can determine if a neighbor node is in the shortest path whenever it received a new data earliest from that node. The sink node will reinforce this shortest path by sending a reinforcement packet with a higher data rate to this neighbor node that forwards it to all the nodes in the shortest path. Other non-optimal paths may be negatively reinforced so that they do not forward data at all or at a lower rate.

Distributed services and applications use the publish and subscribe application programming interface (API) provided by directed diffusion. Through the subscribe() function, an application declares an interest that consists of a list of attribute-value pairs. The subscription is then diffused through the sensor network. A source node may indicate the type of data it offers through the publish() function. It then sends the actual data through the handle returned from the publish() function. The sink node then received the data that has propagated through the sensor network using a recv() function call with the handle returned from the subscribe() call.

4. Distributed Services Overview

Self-organizing sensor applications may be implemented using distributed services that permit the distributed sensor applications to adapt their functionalities and reconfigure their task requirements, manage mobility, and recover from sensor failure. Three main types of distributed services are look-up service, composition service, and adaptation service (see *Figure 1*). These servers enable sensor nodes to form spontaneous communities in ad-hoc sensor networks that may be dynamically reconfigured and hierarchically composed to adapt to real-time information changes and events.

These distributed servers may be replicated for higher availability, efficiency, and robustness. Distributed servers coordinate with each other to perform decentralized services—e.g., distributed look-up servers may work together to discover the location of a particular remote service requested by a node.

4.1. Reconfigurable Smart Components

By exploiting these distributed services, sensor applications can be enabled to be self-aware, self-reconfigurable, and autonomous. Sensor applications running in each sensor node are known as reconfigurable smart components. They are the building blocks for constructing large, scalable, and self-organizing sensor networks. (In this paper, we refer to reconfigurable smart sensor components as smart components or simply sensor nodes.) Smart components may represent sensor nodes or other types of mobile nodes, fixed nodes, or clusters of these nodes. They may simultaneously be service providers for other smart nodes and clients of services that other smart nodes provide.

Smart nodes interact with other smart nodes through well-defined interfaces (for networking and systems operations) that also maintain interaction states to allow nodes to be dynamically reconfigured. These explicit interaction states and behavior information allow localized algorithms with the adaptation servers to maintain consistency when autonomous nodes and clusters are reconfigured dynamically, move around, or recover from failure.

4.2. Look-Up Server

The look-up service enables new system and network services to be registered and made available to other sensor nodes. Methods for calling the services remotely are also provided.

A sensor node that provides a service is called a service provider, and a node that uses the service is called a service client. Because service providers may be introduced or removed from the sensor network at any time, a look-up server is needed to keep track of the availability of these services. A sensor node may register a resource that it maintains or service that it can perform with a look-up server (see *Figure 2*). A look-up server may contain information on services or resources at multiple clusters. Other nodes that require the service may request the service through a look-up server. If the service is recorded in the look-up server, it will return the location of that service to the requesting node. Otherwise, if the service is not recorded in the look-up server of the region, a discovery protocol is invoked to locate the service through other look-up servers (see *Figure 2*). A request message is propagated to all of the look-up servers, and the server that contains the service registration information will return the reply with the service location.

The look-up server that made the request will then store that service location information in its local registration table. At a regular frequency, service and resource registration information may be exchanged among the look-up servers.

4.3. Composition Server

The composition service allows clusters of sensor nodes to be formed and managed. It allows various smart nodes to be

FIGURE 2

Discovery of Services with Distributed Look-Up Server

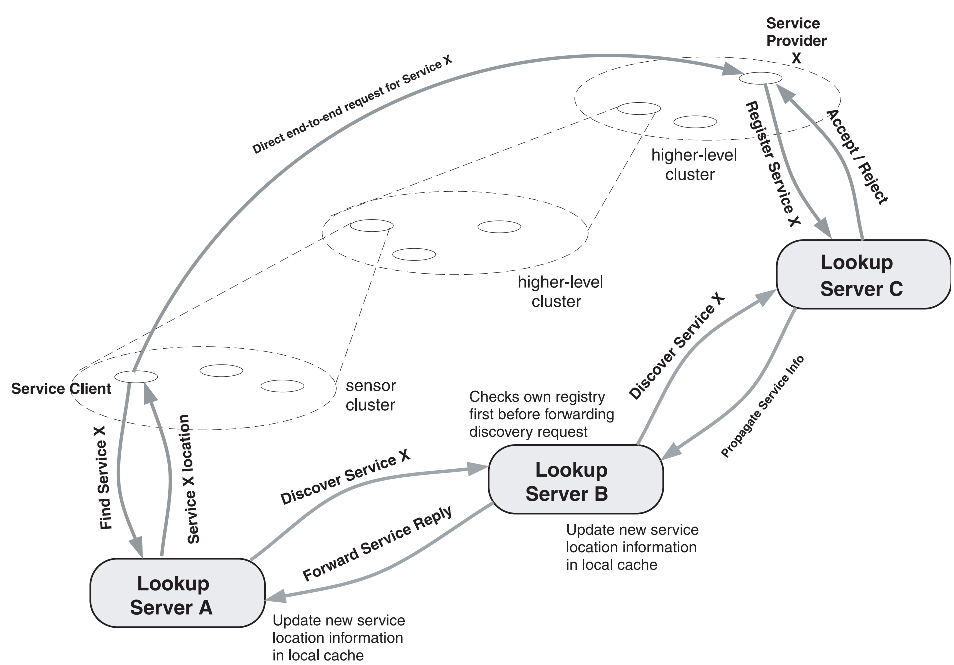

added to (or removed from) clusters in the sensor network. Hierarchical clusters are also possible for larger-scale sensor networks. A cluster of sensors may also provide distributed services by coordinating the tasks among the sensors, such as aggregating summary information. Clustered smart nodes encapsulate the networking and system capabilities provided cooperatively by the group of smart nodes. There will be a head smart node in the cluster that is responsible for the control of the cluster and inter-cluster communications and networking functions. Group communication to nodes in a cluster can be efficiently implemented by sending a message first to the cluster head, which then multicasts it to the member nodes. Member nodes will elect a cluster head from the set of nodes with the most powerful networking and system capabilities. Smart nodes in a cluster may cooperate to perform the networking and system functions for the cluster. Synchronization constraints associated with network protocols and system services among smart nodes may be specified in clustered smart nodes. The capability to specify hierarchical composite clusters enables designers to build large and complex sensor networks by clustering together smaller network-enabled sensor devices at each level.

4.4. Adaptation Server

The adaptation service allows sensor nodes and clusters to reconfigure dynamically as a result of sensor-node mobility, failure, and spontaneous deployment. Adaptation servers utilize information from the composition server, look-up server, and analytical tools to control smart nodes during dynamic reconfiguration and failure recovery of the agile sensor network. Each smart node may execute autonomously to control different network operations in the sensor network and may interact and coordinate independently with other smart nodes to perform collaborative networking operations.

Adaptation servers monitor clusters of smart nodes during normal execution, either through the spontaneous signal from the sensors, probing of the smart nodes, or explicit network-management directives for reconfiguration and failure recovery. When a runtime reconfiguration is requested or triggered, the adaptation server will generate the appropriate schedule of reconfiguration operations that will ensure that the reconfigured and affected sensor nodes are globally consistent. When smart nodes are added or removed from the agile sensor network, a suite of analytical tools may be utilized to ensure that the sensor network still maintains its safety and liveness properties [13].

5. API for Look-Up Service

Applications use the look-up service through the following API. We focus primarily on the look-up service API in this paper, because it is responsible for enabling sensor networks to be self-organizing. In the next section, we will describe how these API functions are used by the various application systems of a surveillance sensor network.

These API functions use the following parameters:

- ***service_type*** is the generic type of service for which there may be several instances. For example, a type of service may be temperature monitoring, whereas specific instances of temperature sensors may be sensor *X* at location *Y*.
- ***service_name*** is the specific name that identifies an instance of a service provider.
- ***input_list*** is a list of attribute-value pairs containing the input parameters to the service invocation.
- ***output_list*** is a list of attribute-value pairs containing the output values from the service invocation.
- ***lifetime*** is the time period in which a service information will be stored in a look-up server.
- ***interface_type*** is one of the following three types of interfaces that the callee of the service uses:
  1. *Location or Address:* This is used by the service client if it knows the interface for interacting with the service provider—e.g., using remote procedure calls. The service client merely needs to retrieve the location or address of the service provider to be used for invoking the service request.
  2. *Interface Definition:* This is used by the service client to retrieve the definition of the interface for interacting with the service provider. The service client must have the interpreter or compiler for the interface definition.
  3. *Mobile Code:* This is used by the service client to retrieve the mobile code that implements the protocol for interacting with the service provider. The mobile code must then be dynamically linked to the service client.

The following is a description of the purposes and the side effects of the look-up service function calls:

- ***service_call(service_name, input_list, output_list, interface_type):*** This function is used by a service client to find and make a call for a service where the service client does not know the location or address of the service provider and/or the interface for using the service. It is implemented as a combination of the look-up_service() and service_exec() calls described below. This function requires specific a service name to be provided. A generic service type cannot be used because several service-provider instances may match the service type.

- ***status = look-up_service(service_type, service_name, input_list, output_list, interface_type):*** This function allows a service client to find the location or address of a service provider and/or the interface for using the service. If service_type is defined and service_name is NULL, then all service providers registered with the look-up server of that type is returned. Service look-up can also be based on cluster or predicate matching. The cluster information and predicate for matching service providers are contained in the input_list. Depending on the interface_type used—i.e., location, interface definition, or mobile code—the respective results of the look-up_service() call will be placed in the output_list.

- ***status = service_exec(service_name, input_list, output_list, interface_type):*** This function allows a node to request for the service and gets the results back from the service provider. The input parameters are supplied by the client via the input_list. The service provider performs the requested service or remote procedure call and returns the results in the output_list. The interface_type defines the method used by the service client to communicate with the service provider.

- ***status = service_register(service_type, service_name, lifetime, input):*** This function allows a service provider to register its service with a look-up server in the region. Services will remain in the look-up server for the lifetime specified by the service provider. In the input_list, the service provider may supply one or more of the following service information: location or address, interface definition, or mobile code. Look-up servers for different regions may coordinate with each other to update their list of service information.

- ***status = service_deregister(service_type, service_name):*** This function allows a service provider to remove its service from the look-up server registry.

6. Implementation of Look-up Service

The look-up service is implemented using a directed diffusion routing protocol that is capable of adapting to changes in the ad-hoc sensor network topology.

6.1 Discovery of Look-Up Servers

When a node (that may implement a look-up server, a service provider, or a service client) moves into a new area or is rebooted, it first sends out a discovery interest to search for a nearby look-up server. If there is a look-up server available, the look-up server will send an acknowledgement with the information of itself—such as its name, ID, location, and other information—to the node by publishing this information.

If several look-up servers reply, the node will accept the first one that replied and ignore the others. The node stores the look-up server information in a table called "look-up table" for future reference.

To reduce bandwidth usage, the service client will only broadcast its discovery interest to within a limited region in its vicinity. The size of the region will depend on number of look-up servers in the sensor network. Each intermediate node runs an application-specific routing filter program that checks whether the node is outside the region for the discovery interest. If it is, the node will not propagate those interests any further. Otherwise, it will broadcast the discovery interest to its other neighbors.

The following are the pseudocodes for look-up server discovery:

Node Discovery Client

```
setupFindingInterest();
while (not found)
    waiting;
unsubscribe();

callback::recv()
{
    if (receive Discovery publication)
       store look-up server info in cache
        set found;
}
```

Look-Up Service Discovery Server Daemon

```
Handle = setupDiscPublication();
while (1)
{
    send(my_location);
}
```

In the node discovery clients, the function setupFindingInterest() sends out the discovery interest. It then waits for the reply from a look-up server. The callback function will be activated when a reply arrives. In the look-up server node, there is a discovery daemon that handles discovery requests and publishes its information.

6.2 Look-Up Server Registration

When a look-up server moves into a new area, it tries to find another existing look-up server and registers with the other look-up server. It sends to other look-up servers in the vicinity the registration interest containing information about itself. Each existing look-up server will update its look-up server table with information from the interest packet and return a list of all look-up servers that it is aware of. It then forwards the registration packet to all those look-up servers, through shortest reinforced routes created by the diffusion protocol.

The following are the pseudocodes for look-up server registration:

Client for Look-Up Service Registration

```
setupRegReqInterest();
while (not succ_reg)
    waiting;
unsubscribe();

callback::recv()
{
    if (recv ACK from Server)
    {
        get shared information;
        set succ_reg;
    }
}
```

Server for Look-Up Service Registration

```
setupSvrRegSvrInterest();
while (not received incoming interest)
    waiting;
// Include shared information;
setupSvrRegACKPublication();
send(ACK, shared LkupSvr information);
unpublish();

callback::recv()
{
    if (matched income interest)
    {
        set signal;
        store in its table();
        // Forward interest to other
        // look-up servers.
        update_lkup();
    }
}
```

The look-up service registration client executing in the look-up server tries to register itself with another look-up server through the look-up service registration server. This registration server will continually publish its information on other look-up servers of which it is aware in response to any registration request from another look-up server. The callback function will be activated when there is a new registration interest. The information from registration interest will be stored in its local table. The information is then forwarded to other look-up servers via the function update_lkup().

6.3 Service-Provider Registration

When a service provider moves into a new area or restarts a service, it first tries to find a new look-up server and registers its service with the look-up server. This is done by sending out the registration interest to the look-up server that includes the information on the service provider itself, such as its location, ID, name, and other information on the service and interface. It then waits for acknowledgement from the look-up server.

For registration that involves transferring mobile codes or interface definitions that implement the service-provider interface, there are three steps in the registration process that includes a Go-Back-N protocol. In the first step, the service provider sends an interest for registering its service. The look-up server then publishes its acceptance or rejection of that request. In the second step, the look-up server sends its interest in receiving the mobile code packets. (Large mobile codes must be fragmented into smaller packets.) The service provider then publishes the packets to the look-up server. When the transmission of the packets is completed, the service provider sends an interest in the third step for receiving an acknowledgement from the look-up server that the mobile code transfer and the registration are successful. If successful, the look-up server will publish an acknowledgement. Otherwise, it publishes a negative acknowledgement.

The following are the pseudocodes for the registration of service providers:

Client for Service Provider Registration (implemented in the service provider)

```
//Step 1.
setupRegInterest();
while (not succ_reg)
    waiting;
unsubscribe();

callback::recv()
{
    if (recv ACK from Server)
    {
        set succ_reg.;
    }
}

//Step 2.
Publish();
Initialize timer;
Send packet with pckcnt and set GBN_Window;

// Step 3.
Create_thread_for(RegACKListener);
Subscribe interest;

Callback::recv()
{
    if (it's a new ACK)
    {
        set GBN_Window;
        set ACKNUM;
        call GBN_timer();
    }
}
```

Server for Service-Provider Registration (implemented in the look-up server)

```
// Step 1.
setupSvrRegSvrInterest();

callback::recv()
{
    StoreInCache();
    Send ACK;
    // Step 2. Interest in receiving
    // mobile code packets
    Subscribe interest();
}

// Step2. Receiving packets
Callback::recv()
{
    StoreInCache();
    // Step 3.
    Send ACK;
}
```

Since diffusion network routing is based on events and notification model of communication, we implemented an end-to-end automatic repeat request protocol based on Go-Back-N to ensure reliable transmission of the mobile codes and interface definition that could be large.

6.4 Service Deregistration

When a service provider wants to terminate its service, it first deregisters the service with the look-up server by sending it a deregistration interest. After the look-up server receives the deregistration interest, it purges the service record from its service table and sends the update interest to all other look-up servers.

6.5 Service Look-Up

When a node needs to find information on a specific service and the relevant service provider, it calls its local service look-up function that sends a look-up interest to the look-up server to fetch the updated service information. After the look-up server received the look-up interest, it searches its local service table first. If it finds the service record, it will return a data packet to the node by publishing the data. Otherwise, it will send the look-up request interest to all other look-up servers. If the service still cannot be found, it will return a failure message to the requesting node.

The requesting node can look up a service using a service type instead of a specific service name. The service type is the generic type of the service—for example, a generic type of a detector may be "Detector", whereas the service name of that detector is "Detector_AA_trace_01." If the look-up server received a service-type interest, it will return the number of matched service providers as requested. The default number of service providers is 1.

The following are the pseudocodes for service look-up:

Client of Service Look-Up (executes in the service requester)

```
GetserviceCache();
if (no record)
{
    setupLKUPInterest();
    while (no succ_sig)
        wait;
    unsubscribe();
}

callback::recv()
{
    if (recv look-up return)
       cache provider's info;
        set succ_sig;
}
```

Server of Service Look-Up

```
setupLKUPSvrInterest( );
while (1)
    ;

callback::recv()
{
    unpack interest;
    look-up the cache;
    if (no record in cache)
    {
        setupLKUPInterest();
        // Forward the interest to
        // other LKUP Servers.
        While(no succ_sig)
            Waiting;
        Unsubscribe();
    }
    store results in local table;
    setuplook-upPublication();
    send();
    unpublish();
}
```

The look-up server sets up a local interest by calling the function setupLKUPSvrInterest(). When it receives a corresponding matched global interest from the client, this local interest will activate the callback function to search and return the service provider's information. It returns this information through publish and send functions. If no information on the service provider is found in the local table, the request interest will be forwarded to other look-up servers. The results will be stored in its local table and then sent to the requesting node.

6.6 Remote Service Execution

When a node wants to execute a service from a service provider, it searches its local service table to get the information on the service and the service provider. If it cannot find the relevant information, it will call the service look-up function to update its local service table. Three types of service interfaces may be specified by the service provider's information: Location or address of the service provider with known interface, interface definition of the service, or mobile code for the interface protocol. These interfaces were described in Section 5.

Consider the example in which the service client uses remote procedure calls (RPC) to request remote services from the service provider. RPCs are implemented as follows: The call to service_exec() will first send an interest to the service provider through subscribe() function in the SetupserviceExecInterest() function. The service provider then sends a data to the client containing the permission to send the request. The client then sends a request and all the input data. If the input data is large, several packets may be sent reliably through an automatic repeat request protocol with retransmission. The service provider will then process the service remotely. The client will request the result of the service through another subscription. The provider then returns the result in a data in response to the interest.

The following are the pseudocodes for remote service execution:

Client of Service Execution

```
SetupserviceExecInterest();
while (not execget)
    wait;
unsubscribe( );

callback::recv()
{
    if (recv Result from service Provider)
    {
        collect service Execution result;
        set succ_reg.;
    }
}
```

Server of Service Execution

```
SetupserviceExecSvrInterest();
while(1)
    sleep;

callback::recv()
{
    if(recv service Execute Interest)
    {
        collecting the data;
        setupserviceExecSvrPublication();
    }
}
```

7. Support for Sensor Node Mobility

Look-up servers support the mobility of sensor nodes. When a sensor node moves to a different cluster at another location, it notifies, whenever possible, the previous look-up server that it is moving. When it arrives at another cluster in a new location, it will register with the new look-up server that will notify the previous look-up server (see *Figure 3*). The new look-up server will propagate the change in the node's location to other look-up servers. The look-up server responsible for a sensor node that is interacting with the mobile node will notify the sensor node of the service location change. Existing interactions between the mobile node and other nodes will thus be handed over to the new location. Adaptation servers may be involved in the handover operation to preserve global consistency during the handoff resulting in uninterrupted use of the service.

FIGURE 3

Mobility of Smart Nodes

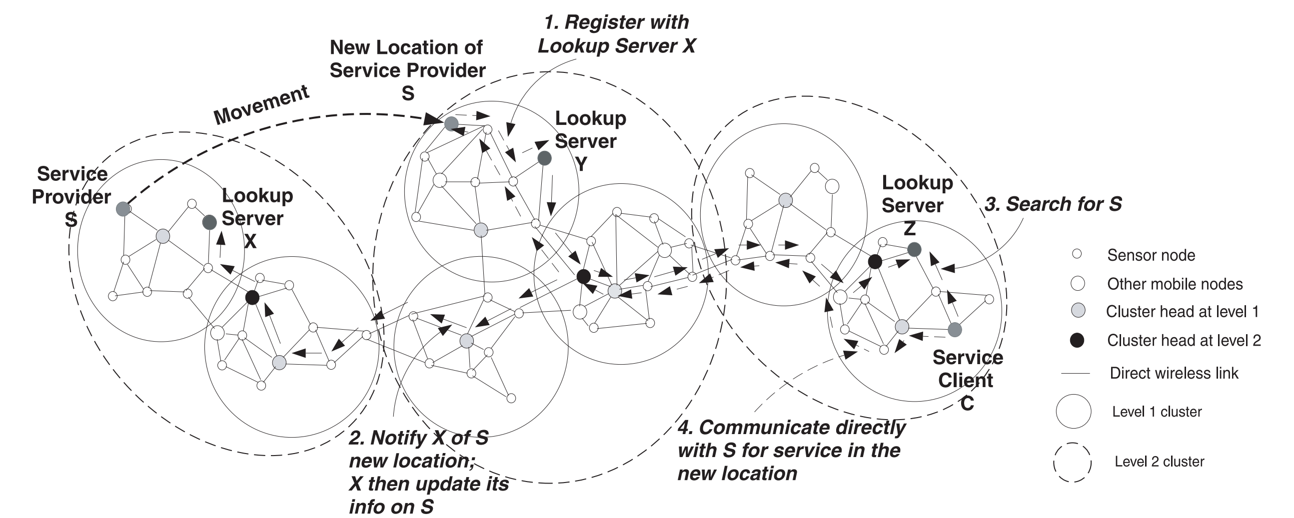

8. Application Systems

In the following, we discuss how the look-up service enables various application services for collaborative signal processing to be self-configurable in ad-hoc sensor networks.

In target recognition and tracking [3], the multiple reading of the sensor values can be statistically combined to derive more accurate tracking data. Sensors in a region may be clustered together through the distributed composition server, where their sensor reading may be combined using a weighted voting algorithm [18] to provide more accurate data. For each cluster, a sensor node may be elected to be the head of the cluster. Readings from multiple sensors in the clusters will be propagated to the cluster head through diffusion routing, where the algorithm is applied. Results from each cluster head may be propagated to higher-level cluster heads for data fusion.

In many sensor applications, data generated by the sensor node may be very large. The cost of propagation of these sensor data as in the previous scheme will be prohibitive. One solution to address this problem is to use mobile agents to migrate from node to node to perform data fusion using the local data [17]. Instead of transferring large amounts of data throughout the sensor network, this approach only transfers the mobile agent code that is smaller than the sensor data. The result is an improvement in the execution time of the collaborative signal-processing algorithm.

A sensor node that needs to execute a mobile agent code must download the mobile code from a mobile code repository manager in its region. It may use the look-up_service() call to first locate the repository manager that stores the relevant mobile code. It then calls service_exec() to the repository manager to download the code. An alternative method is to store the mobile code with the look-up server. The sensor node can then use look-up_service() to retrieve the mobile code directly from the look-up server.

The entire surveillance area is split into several sub-areas. Each sub-area is controlled by a signal-processing agent, also known as a processing element (PE), which may dispatch several mobile agents into that sub-area. The signal processing agent may register itself with the look-up server using service_register() to allow other nodes, such as the mediators, to access its tracking results. Each mobile agent will migrate from node to node to perform data fusion, for instance using a multi-resolution data integration algorithm [17]. Each mobile agent, addressed by an identification, contains an itinerary, data, method, and interface. The identification is a 2-tuple composed of the identification of the dispatcher and the serial number assigned by the dispatcher. The itinerary describes the migration route assigned by the dispatcher. The data is the agent's private data that contains the integration results. The method describes the multi-resolution data integration algorithm. The interface is the function by which the agent communicates with the PE and for the PE to access the agent's private data.

When mobile agents migrate from node to node, the results of the sensor integration algorithm from previous nodes are cached in the mobile agents. This state information must be transferred with the mobile agent as it migrates from node to node. Agent migration and state transfer are supported by the distributed adaptation service. As the mobile agents visit each node, it may register with the look-up server. The primary signal-processing agent (PE), may find these mobile agents as they move around and retrieve results using the service_call() function. A mobile agent may also retrieve intermediate results from other mobile agents through service_call().

For migration of mobile agents far beyond a region managed by a look-up server, the agent will register with another look-up server in the new region. A service client may try to contact the mobile agent through the first look-up server and may need to use several intermediate look-up servers to get the information about the mobile agent's current location.

Multi-resolution signal processing algorithms may be implemented using a hierarchical clustering of sensors provided by the distributed composition server. Results from a sensor cluster may be passed to agents responsible for signal processing for higher-level clusters. Agents for higher-level clusters may send messages to multiple sensor clusters using a group communication method provided by the service_exec() call by specifying the appropriate service_type parameter.

9. Related Work

In recent years, the development of new integrated low-power sensor devices [11,19,21] has spurred more research on sensor networks for remote surveillance [3], pervasive computing, mobile computing, and computing continuum. Some research on sensors and RFID tags have also focused on the technical issues in wireless links between the sensor and the interrogator [12], low power consumption [19], sensor types and capability [14], position determination [20], and identification. Many of the current commercial systems transmit data from sensors to interrogators through wireless links, but forward them through direct static connections to a fixed network, either through a cable or wireless link. However, over-reliance on fixed network infrastructure is not feasible for a large sensor information network. Instead, sensors must themselves be capable of assembling ad-hoc networks when they are deployed spontaneously in an area. The sensors must self-organize in spite of changes and failure in the sensors and the network topology.

Several novel concepts and protocols have been developed for many aspects of self-organizing sensor networks' design and implementation. In ad-hoc network routing, localized algorithms have been developed for autonomous sensor devices to route information in an energy efficient way. A directed diffusion routing protocol [8] based on the localized computation model provides energy-efficient and robust communication for dynamic network with small incremental changes. For dynamic networks with large-scale changes and high level of mobility, directed diffusion may not adapt very well. Similar diffusion routing concepts have also been presented in other work [16]. Another localized protocol for information dissemination in sensor networks uses meta-data negotiation to eliminate redundant transmissions [10]. Other self-organizing network routing protocols are dynamic source routing [9] and destination-sequenced distance vector [15], although it is not clear if these algorithms are energy efficient enough for sensor networks. In this paper, we allow sensors to form high-level clusters and use directed diffusion within clusters. Clusters may be formed using a localized algorithm [6] for coordinating among sensors to elect extremal sensors.

The discovery of services in mobile systems provides critical support for self-organizing sensor systems when sensors are being deployed and removed on the fly. In Jini [1], service discovery relies on mobile Java codes and is implemented based on transmission control protocol (TCP) and user datagram protocol (UDP). It is not clear how these may be implemented using data-centric, ad-hoc sensor networks with services based on more generic mobile codes. Service-location protocol (SLP) [7] is an Internet Engineering Task Force (IETF) protocol for service discovery that is designed solely for IP–based networks. Bluetooth [2] devices have a range of 10 meters and can directly communicate with at most seven other Bluetooth devices in a piconet. Bluetooth service-discovery protocol (SDP) [2] allows devices to browse and retrieve services by matching service classes or device attributes. Only services within the range of the device are returned. Our look-up service may retrieve services that could be multiple hops from the requesting node.

Communications in sensor network are data-centric, since the identities of the numerous sensors are not as important as the data that they contain [5,6]. The networking infrastructure may provide a more efficient dissemination of data through replication, caching, and discovery protocols [5]. Communications protocols must be energy efficient because sensors [11] have very limited energy supply. In our architecture, the caching and aggregation of information may be provided by some sensors. The discovery of these sensors can be made through the distributed look-up servers that are implemented using diffusion routing. Changes in the sensor network are propagated to other caching and aggregation services through the adaptation servers. Our framework facilitates the consistent adaptation of networking and system services as well as distributed sensor applications. Unlike other centralized control networks [4], our servers are associated only with sensors in a vicinity. Servers in different clusters will coordinate among themselves through information diffusion.

10. Conclusions

We have described how distributed look-up services can be implemented using directed diffusion network protocols. The look-up service supports self-organization of sensor networks that are useful for many sensor applications, such as sensor information retrieval and remote surveillance. The look-up service is implemented in the overall distributed service architecture that consists of three basic distributed servers: look-up servers, composition servers, and adaptation servers. Through the look-up servers, sensor nodes may be placed together impromptu in spontaneous environments, and these sensor nodes will immediately know about the capabilities and functions of other sensor nodes and work together as a community system to perform cooperative tasks and networking functionalities. These distributed services are implemented using directed diffusion network routing that provides energy-efficient and data-centric data communications. While diffusion routing can adapt dynamically to limited mobility and topological change, the distributed look-up services support large mobility and changes in sensor nodes in a large-scale ad-hoc network. Diffusion uses a data-centric communications model, whereas in distributed sensor applications, it may be more convenient to use end-to-end process-oriented communications.

References

[1] K. Arnold, et al., *The Jini Specification*. Addison Wesley, 1999.
[2] *Specification of the Bluetooth System*, available at http://www.bluetooth.com/developer/specification/specification.asp
[3] R. R. Brooks, "Chapter 26: Modern Sensor Networks," *CRC Handbook of Sensor Fusion*, accepted for publication.
[4] Echelon, The LonWorks Company. LonWorks Solutions. http://www.echelon.com/Solutions/
[5] M. Esler, et al., "Next Century Challenges: Data-Centric Networking for Invisible Computing," *ACM Mobicom*, 1999.

[6] D. Estrin, et al., "Next Century Challenges: Scalable Coordination in Sensor Networks," *ACM Mobicom*, 1999.

[7] E. Guttman, "Service Location Protocol: Automatic Discovery of IP Network Services," *IEEE Internet Computing*, Vol. 3, No. 4, July/Aug. 1999, pp. 71–80.

[8] C. Intanagonwiwat, R. Govindan, D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *ACM Mobicom*, 2000.

[9] D.B. Johnson and D.A. Maltz, "Dynamic Source Routing in Ad-Hoc Wireless Networks," in T. Imielinski and H. Korth, editors, *Mobile Computing*, pages 153-181, Kluwer Academic Publishers, 1996.

[10] J. Kulik, et al., "Negotiation-based Protocols for Disseminating Information in Wireless Sensor Networks," *ACM Mobicom*, 1999.

[11] J.M. Kahn, et al., "Next Century Challenges: Mobile Networking for Smart Dust," *ACM Mobicom*, 1999.

[12] A. Lim and K. Mok, "Wireless Media Access Control for Highly Mobile Information Servers: Simulation and Performance Evaluation," *ACM Mobile Computing and Communication Review*, Vol 1, No. 2, 1997.

[13] A. Lim, "Automatic Analytical Tools for Reliability and Dynamic Adaptation of Complex Distributed Systems," *IEEE ICECCS*, Nov 1995.

[14] H. T. Nagle, et al., "The How and Why of Electronic Noses," *IEEE Spectrum*, September 1998, pp. 22-34.

[15] C. E. Perkins and P. Bhagwat, "Routing Over Multi-Hop Wireless Network of Mobile Computers," in *Mobile Computing*, edited by T. Imielinski and H. Korth, Kluwer Academic Publishers, 1996, pp. 183–206.

[16] R. Poor, "Hyphos: A Self-Organizing, Wireless Network," Master's Thesis, MIT Media Lab., June 1997.

[17] H. Qi, S. S. Iyengar and K. Chakrabarty, "Distributed Multi-Resolution Data Integration Using Mobile Agents," *Proc. IEEE Aerospace Conference*, 2001.

[18] D. G. Saari, "Geometry of Voting: A Unifying Perspective," Proc. *Workshop on Foundations of Information/Decision Fusion with Applications to Engineering Problems*, DOE/ONR/NSF, Washington DC, August 1996.

[19] The Ultra Low Power Wireless Sensors project. http://www-mtl.mit.edu/~jimg/project_top.html

[20] J. Werb, et.al., "Designing a Positioning System for Finding Things and People Indoors," *IEEE Spectrum*, Sep 98, pp. 71-78.

[21] G.J. Pottie and W. J. Kaiser, "Wireless Integrated Network Sensors," *Communications of the ACM*, Vol. 43, No. 5, May 2000.

Notes

1. This research is supported in part by the DARPA Sensit Program and Space and Naval Warfare Systems Center, Department of Navy, under Grant N66001-00-1-8948.

Wireless LAN: Its Benefits and Implications

Gordon K. Lin
President and Chief Executive Officer
Lintel Communications, Inc.

Introduction

Wireless local-area network (LAN) technology is rapidly becoming a crucial component of computer networks. According to the IEEE 802.11 wireless LAN standard, wireless network technology has emerged from the world of proprietary implementations to become an open solution for providing mobility as well as essential network services where wireline installations proved impractical. Now companies and organizations are investing in wireless networks at a higher rate to take advantage of mobile, real-time access to information.

Wireless LAN Benefits

The emergence and continual growth of wireless LANs are being driven by the need to lower the costs associated with network infrastructures and to support mobile networking applications that offer gains in process efficiency, accuracy, and lower business costs. This section explains the mobility and cost-saving benefits of wireless LANs.

Mobility

Mobility enables users to move physically while using an equipment/instrument, such as a handheld personal computer (PC) or data collector. Many jobs require workers to be mobile, such as inventory clerks, health-care workers, policemen, and emergency-care specialists. Of course, wireline networks require a physical connection between the user's workstation and the network's resources, which makes access to these resources impossible while roaming about the building or elsewhere. This freedom of movement results in significant return on investments (ROIs) due to gains in efficiency.

Mobile applications requiring wireless networking include those that depend on real-time access to data—usually stored in centralized databases. If your application requires mobile users to be aware immediately of changes made to data, or if information put into the system must immediately be available to others, then you have a definite need for wireless networking. For example, for accurate and efficient price markdowns, many retail stores use wireless networks to interconnect handheld bar-code scanners and printers to databases having current price information. This enables the printing of the correct prices on the items, making both the customer and the business owner more satisfied.

Installation in Difficult-to-Wire Areas

The implementation of wireless LANs offers many tangible cost savings when performing installations in difficult-to-wire areas. If rivers, highways, or other obstacles separate buildings that you want to connect, a wireless solution may be much more economical than installing physical cable or leasing communications circuits, such as T1 service or 56K lines. Some companies spend thousands or even millions of dollars to install physical links with nearby facilities. If you are facing this type of installation, consider wireless networking as an alternative.

The asbestos found in older facilities is another problem that many companies encounter. The inhalation of asbestos particles is extremely hazardous to your health; therefore, you must take great care when installing network cabling within these areas. Some companies, for example, remove the asbestos, making it safe to install cabling. This process is very expensive. Obviously, the advantage of a wireless LAN in asbestos-contaminated buildings is that you can avoid the asbestos removal process, resulting in tremendous cost savings.

Some municipalities may restrict you from permanently modifying older facilities with historical value. This could limit the drilling of holes in walls during the installation of network cabling and jacks. In that case, a wireless network might be the only solution. Right-of-way restrictions within cities and counties may also block the digging of trenches in the ground to lay optical fiber for networked sites. Again, in this situation, a wireless network might be the best alternative.

Increased Reliability

A problem inherent to wired networks is downtime due to cable faults. In fact, a cable fault is often the primary cause of system downtime. Moisture erodes metallic conductors via water intrusion during storms and accidental spillage or leakage of liquids. With wired networks, a user might accidentally break his/her network connector when trying to disconnect his/her PC from the network to move it to a different location. Imperfect cable splices can cause signal reflections that result in unexplainable errors. The accidental cutting of cables can bring down a network immediately. These problems interfere with users' ability to use network resources, causing destruction for network managers. An advantage of wireless networking, therefore, results from the use of less cable. This reduces the downtime of the network and the costs associated with replacing cables.

Reduced Network-Installation Time

The installation of cabling is often a time-consuming activity. For LANs, installers must pull twisted-pair wires or optical fiber above the ceiling and drop cables through walls to network jacks that they must affix to the wall. These tasks can take days or weeks, depending on the size of the installation. The installation of optical fiber between buildings within the same geographical area consists of digging trenches to lay the fiber or pulling the fiber through an existing conduit. You might need weeks or possibly months to receive right-of-way approvals and dig through ground and asphalt.

The deployment of wireless LANs greatly reduces the need for cable installation, making the network available for use much sooner. Thus, many countries lacking a network infrastructure have turned to wireless networking as a method of providing connectivity among computers without the expense and time associated with installing physical media. This is also necessary within the United States to set up temporary offices and rewire renovated facilities.

Long-Term Cost Savings

Companies reorganize, resulting in the movement of people, new floor plans, office partitions, and other renovations. These changes often require recabling the network, incurring both labor and material costs. In some cases, the recabling costs of organizational changes are substantial, especially with large enterprise networks. The advantage of wireless networking is again based on the absence of cable: You can move the network connection by simply relocating an employee's PC.

Wireless LAN Implications

Wireless LANs offer tremendous benefits, as described in previous section of this paper. However, project managers and design engineers should take the following issues into consideration when implementing wireless networking:

- Multipath Propagation
- Path Loss
- Radio Signal Interference
- Network Security
- Application Connectivity
- Installation Predictability

Multipath Propagation

Transmitted signals can combine with reflected ones to corrupt the signal detected by the receiver. This is known as multipath propagation. Delay spread is the amount of delay experienced by the reflected signals compared to the primary signal. As delay spread increases, the signal at the receiver becomes more distorted and possibly undetectable even when the transmitter and receiver are within close range. Office furniture, walls, and machinery are obstacles that can redirect parts of the transmitted signal. Wireless LAN manufacturers compensate for the effects of multipath propagation by using special processing techniques. As examples, equalization and antenna diversity are methods for reducing the number of problems arising from multipath propagation.

Path Loss

Path loss between the transmitter and receiver is a key consideration when designing a wireless LAN solution. Expected levels of path loss, based on the range between the transmitter and receiver, provide valuable information when determining requirements for transmit power levels, receiver sensitivity, and signal-to-noise ratio (SNR). Actual path loss depends on the transmit frequency, and it grows exponentially as the distance increases between the transmitter and receiver. With typical indoor applications, the path loss increases approximately 20 decibels (dB) every 100 feet.

Radio-Signal Interference

Wireless networks can interfere with other nearby wireless networks and radio wave equipments. Radio-signal interference can be inward or outward. A wireless LAN can experience inward interference from the harmonics of transmission system or other products (e.g., microwave ovens) using similar radio frequencies in the local area. Outward interference occurs when a wireless network's signal disrupts other systems, such as adjacent wireless LANs and navigation equipment on aircraft.

When dealing with interference, you should coordinate the operation of radio-based wireless network products with your company's frequency-management organization, if one exists. Government organizations and most hospitals generally have people who manage the use of transmitting devices. This coordination will avoid potential interference problems. If no frequency-management organization exists within your company, run some tests to determine the propagation patterns (radio frequency [RF] site survey) within your building. These tests let you know whether existing systems might interfere with, and thus block and cause delay to, your network. You will also discover whether your signal will disturb other systems.

Network Security

Network security refers to the protection of information and resources from loss, corruption, and improper use. The main security issue with wireless networks is that they intentionally propagate data over an area that may exceed the limits of the area that the organization physically controls. For instance, radio waves easily penetrate building walls and are receivable from the facility's parking lot and possibly a few blocks away. Someone can passively retrieve your company's sensitive information from this distance without being noticed. Someone can maliciously jam the radio-based network and keep you from using the network. Remember, most wireless networks utilize a carrier sense protocol to share the use of the common medium. If one station is transmitting, all others must wait. Someone can easily jam your wireless network by using a wireless product of the same manufacturer that you have within your network and setting up a station to resend packets continually. These transmissions block all stations in that area from transmitting, thereby making the network inoperable.

Wireless LAN vendors solve most security problems by restricting access to the data. Most products require you to establish a network access code and set the code within each workstation. A wireless station will not process the data unless its code is set to the same number as the network. Some vendors also offer encryption as an option.

Application Connectivity
The use of traditional wired-based protocols over wireless networks introduces problems with maintaining connections between the user's device and the application residing on a server. Transmission control protocol (TCP)/Internet protocol (IP), for example, over wireless networks is susceptible to losing connections, especially when the device is operating in an area with marginal wireless network coverage.

A solution to this problem is to use wireless middleware software, which provides intermediate communications between the end-user devices and the application software located on a host or server. The middleware enables highly efficient and reliable communications over the wireless network, while maintaining appropriate connections to application software and databases on the server/host via the more reliable wired LAN.

The mobile nature of wireless networks can offer addressing problems as well. Most networks require the IP address loaded in the user's device to be within a specific address range to maintain proper connections with applications. When a user roams from one IP subnet to another with a wireless device, the device and the application may lose the capability to connect with each other. As a result, implementers should consider the use of MobileIP as a mean of maintaining connectivity while traversing different IP domains.

Installation Predictability
With wired networks, planning the installation of cabling is fairly straightforward. You can survey the site and look for routes where installers can run the cable. Once the design is complete, installers can run the cables, and the cable plant will most likely support the transmission of data as planned. A wireless LAN installation is not as predictable. It is difficult to design the wireless network by merely inspecting the facility. Predicting the way in which the contour of the building will affect the propagation of radio waves is difficult. Walls, ceilings, and other obstacles attenuate the signals more in one direction than the other and even cause some waves to change their paths of transmission. Even the opening of a bathroom door can change the propagation pattern.

To avoid installation problems, an organization should perform propagation tests to access the coverage of the network. Neglecting to do so may leave users outside the propagation area of wireless servers and access points. Propagation tests give you the information necessary to plan wired connections between access points, allowing coverage over applicable areas.

Conclusion

Wireless LAN is moving toward open architecture implementation due to the introduction of the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standard and is growing rapidly. It offers great benefits such as mobility and cost savings. When implementing wireless networking, the following issues should be taken into consideration: multipath propagation, expected path loss, radio-signal interference, network security, connectivity, and installation predictability.

3G Internet Mobility

Bennesa Lyon
Senior Business-Development Manager, Mobile Internet
Compaq Telecom

Brian Haley
Principal Engineer/Mobile IPv6 Technical Leader
Compaq Tru64 UNIX Networking Group

Jamey Hicks
Principal Member of the Technical Staff
Compaq Research

Yanick Pouffary
Senior Member of the Technical Staff
Compaq EMEA

Mobility—it is the hallmark of 21st century communications, with mobile voice services having already transformed both business and personal interactions. But the full potential of mobile communications won't be realized until the capability extends to data. This next stage—extending seamless mobility to the Internet—has begun, and it will represent the true mobile communications revolution, empowering users with anywhere, anytime, anyway access to the world's greatest reservoir of information and entertainment, and the capability to instantly share this wealth with others.

The development of third-generation (3G) networks, with their tremendous capabilities, will spur the creation of unprecedented amounts and types of content available on the Internet, including video, voice, and images. Combine that with the number of wireless subscribers expected to reach one billion by 2004[1] and the number of Internet users projected to reach 600 million a year later[2], and the demand for mobile Internet access is inevitable. So, too, will be the demand for convergent devices that provide both voice and data.

The communications industry already provides rudimentary mobile data services via wireless application protocol (WAP) and general packet radio service (GPRS) devices. On the horizon, though, are an exciting wave of mobile communications products and services made possible by a revolutionary new network-layer protocol called Internet protocol version 6 (IPv6). For the communications industry, mobile IP will indeed mean going new places. And Compaq Telecommunications will be there with the cutting-edge products and technology that make it possible.

Mobile IP: Telecom Challenges and Opportunities

No telecom company needs to be convinced about the desirability or viability of mobile communications. In a very short time, mobile telephony has changed the world. Soon, the number of mobile handsets will overtake the number of fixed lines. For business, mobile communications today means that deals are hastened, products and services are more rapidly developed and delivered, and productivity is enhanced. On a personal level, mobile communications today provides the comfort and assurance that friends and family are immediately available.

The coming ascendance of mobility in the computing world will be just as revolutionary. Soon enough, we will live in a world where accessing e-mail, receiving sports scores or stock prices, playing games, or downloading critical information will be as easy with a data device as it is with a cellular one. Indeed, *these devices will be one and the same*.

After all, consumers will ask, why should mobile communications exclude the world of multimedia content that is data? Why not download information from someone and immediately discuss it with them on the same device, regardless of where you are or where they are, or where either of you is going?

Compaq believes—and the historical development of the Internet confirms—that consumers will not just want more and better products and services, but that they'll also want to access these benefits as conveniently as possible. Until now in the computing world, that meant tailoring information and devices. But as important as those steps have been, the next advance is even more impressive: making the instant information of the Internet available in a way that addresses today's mobile work and lifestyle demands.

For telecom companies, the opportunities are immense. One obvious advantage is the sheer traffic volume that will be enabled by mobile Internet protocol (IP). But mobile IP will also make possible an astonishing new range of services, from gaming to global messaging services to home gateways that control your air conditioning or security. These new services will open the door for usage-based billing and other revenue streams for service providers.

Even at this early stage, it's clear that consumers want these capabilities.

This new world of mobile communications, where voice and data join, will also challenge telecom companies. It will first and foremost require a mindset that recognizes the importance and viability of mobility as it pertains to the Internet. The rewards of mobile IP won't be realized unless and until the new networking protocol known as IPv6 is embraced by the overall ecosystem of communications—the Internet, cellular, land-line, transport, and access networks

Mobile IP will take all of us one step closer to realizing the true meaning and potential of mobile communications. Consumers and the businesses that serve them will reap the rewards of convergence, as telephony and network computing merge. For telecommunications companies, Internet service providers (ISPs), and content developers, the time to prepare for this new world is *now*.

Mobile Internet Access Today: Limits and Frustrations

Back in the days when the Internet was a data-sharing network for the military and a few research facilities, no one foresaw its growth as a communications powerhouse, and mobility was barely a concept. The current- IP, Internet protocol version 4 (IPv4), has been remarkably useful, as the growth of the Internet attests, but its limitations are now obvious.

The most pressing deficiency is a limited number of addresses allowed by IPv4's 32-bit address. Fortunately, the Internet Engineering Task Force (IETF) has developed IPv6—a next-generation protocol with a 128-bit address that not only resolves the address shortage and other issues (see Appendix), but also provides the foundation for greatly improved mobile access.

Because mobility was not a factor when IPv4 was developed, there was no need to distinguish between who you were and where you were; it was assumed they were one and the same. There is no information in the IPv4 address that indicates a new geographic point of attachment.

As a result, today's users have portable networking capability, but it's far from ideal. The correct delivery of packets requires a new IP address associated with the new point of attachment. One has to establish a new connection whether one moves 100 feet or 1,000 miles. As far too many people know, connecting to the Internet from different locations can be a very frustrating exercise.

It can be a costly exercise as well. Because mobility via IPv4 necessitates informing any agent in the routing process about a new location, additional infrastructure is required—something not always deployed in IPv4 nodes. A system has to be configured with a new address, a correct netmask, and a new default router every time one establishes a connection from a new location. Consequently, the time and effort benefits one might realize from having mobile Internet access can be dissipated because of the hardships required to attain that very access—hardly what today's time-pressed consumers expect or demand.

In designing IPv6, IETF added key functionality to the protocol so that mobile IPv6 could take advantage of such features as next headers, destination options, source routing, and neighbor discovery. Then the IETF was able to require key mandatory support elements in all IPv6 implementations to assist with those features.

Mobile IPv6: Reaching the Net Anywhere, Anytime, Anyway

With mobile IPv6, location is no longer an issue when connecting to the Internet. The mobile node sends information about its point of attachment to a home agent—a node on the home network that allows the mobile node to be reachable at its home address, regardless of its actual geographic location. Packets addressed to the mobile node are intercepted by the home agent and tunneled directly to the mobile node's current location.

Mobile IPv6 resolves the mobility issue by utilizing two addresses for a mobile node: a home address and a care-of address. The care-of address is created whenever the mobile node changes its point of attachment to the Internet and is no longer connected to its home network.

The care-of address can be assigned by a dynamic host configuration protocol (DHCP) server, but one of the great advantages of IPv6 is stateless auto-configuration, where the mobile node creates a unique IPv6 address by combining its local-area network (LAN) media access control (MAC) address with a prefix provided by the network router. There's no need for a manually configured server because no server has to approve or distribute an address. This reduces end-user costs because trained staff is no longer required at the receiving network for the mobile node.

Whenever the mobile node moves, its new care-of address is first registered with the home agent in a process known as "binding." It does this by sending a packet, called a binding update, that contains the mobile node's care-of address and registration lifetime, which dictates the length of time to use the binding, along with a home-address destination option that contains the mobile node's home address. For authentication purposes, the home agent registers the binding and returns a binding acknowledgement to the mobile node.

With mobile IPv6, the home agent automatically forwards data sent to the mobile node's home address to the care-of address in a process known as tunneling. Mobile IPv6 also allows for route optimization, whereby data is sent directly from a correspondent node—a mobile or stationary node that directly communicates with the mobile node, utilizing such applications as e-mail, instant messaging, and streaming audio or video—to the mobile node. The mobile node accomplishes this by sending a binding update to inform the correspondent that it is no longer at home. The correspondent node then sends packets directly to the care-of address of the mobile node, with a routing header that contains the mobile node's home address.

When the mobile node sends packets to another node, its care-of address is set as the source address, and a home-address destination option is included. This preserves the end-to-end model of transmission control protocol

(TCP)/IP, allowing mobility without interrupting any open connections.

For mobile users, location will no longer be an obstacle when connecting to the Internet. The additional time, infrastructure, and cost of today's mobile connection will no longer exist—factors that will certainly spur greater use of these connections.

The Mobile IP Timetable

Mobile IP will one day certainly become as commonplace as mobile telephony, but it won't happen overnight. To talk about the future of mobile IP is to talk about the future of IPv6.

For some time, we will live in a dual IP world with IPv4 and IPv6 coexisting. Fortunately, IPv6 was designed with such a transition in mind. The IETF recognizes that upgraded hosts and routers will need to retain downward compatibility with IPv4 devices for many years. It's expected that we'll live in this dual-IP world for at least another decade.

The global benefits of IPv6—not just mobility, but improved security, functionality, and quality of service (QoS)—will only be realized when IPv6 is the dominant protocol. In essence, the weakest link determines the strength of the network, but many factors will work in favor of IPv6's ascendance. IPv6 has already been adopted as the protocol for next-generation networks by the third-generation partnership project (3GPP), a worldwide standards-setting organization. This endorsement will greatly speed the adoption of IPv6 by future networks, beginning with 3G mobile networks.

Underlying factors also favor IPv6. The lack of addresses in developing and emerging economies; the increasing popularity of wireless in those lands and others; the continued rise of the Internet as the single most robust source of information; the need for anywhere, anytime, anyway access; the natural progression toward convergence—all these point to a bright future for IPv6.

Compaq expects IPv6 deployment to occur in stages, beginning with regional IPv6 networks. These networks will be deployed in order to launch new subscriber services and to gain competitive advantages. These regional IPv6 networks will interoperate with national and international IPv4 networks, thereby assuring global, end-to-end service delivery. Subsequent stages will feature the development and deployment of additional products and services made possible by the advanced features of the new infrastructure. In time, the entire national and global IPv4 infrastructure will migrate to the superior IPv6.

As the 3G infrastructure is built around the world, IPv6 is certain to be a core component. Just as surely, the vast improvements and capabilities of 3G networks will spur a multitude of innovative services accessed by an equally wide array of devices. Making it all possible will be IPv6.

Appendix: Other Advantages of IPv6

While IPv6 makes mobile Internet connections possible, that's only one of its benefits. It's not simply a new address; it provides for a "smarter packet" as well. IPv6 provides increased functionality in such commercially important areas as security and QoS and opens the door to new services.

Security—a crucial element for modern e-commerce—was not a priority when IPv4 was developed. With IPv4, servers usually can't determine whether packets are being received from a legitimate end node. Source address masquerading, known as spoofing, can be used to acquire confidential data or to gain control of a server. Solutions such as firewalls can combat these challenges, but they often hinder connectivity as well.

With IPv6, necessary authentication, security encryption, and data integrity safeguards are an integral part of the protocol. The IPv6 standards-based authentication header extension assures that a packet is truly from its source address. End-to-end encryption at the network layer is provided by another standard header extension—the packet itself is never touched.

Quality of service (QoS) issues are also addressed by IPv6. This improved protocol enables differentiation between non-urgent communications and highly critical applications, such as video conferencing. While these capabilities can be constructed within the IPv4 framework, they're built-in with IPv6. A new traffic flow identification field lays the foundation for QoS functions as bandwidth reservation. In addition, traffic flows can be distinguished for best routing, and those labels can be set to assure a desired security level or cost. This capability opens the door to new services—and increased customer satisfaction as well.

IPv6 also advances the art of multicasting, which is important because there will be increased demand for streaming audio, video, and animated content. The improved address defines a large multicast address space, thereby limiting the degree to which multicast routing information is carried throughout an enterprise. IPv6 hosts and routers are required to support multicasting. And IPv6 also introduces the concept of "anycast" services, whereby a group of nodes can be designated as an anycast group, and a packet addressed to the group's address is delivered to only one of the nodes.

Notes

1. *Mobile Computing Forecast: 2000*, Gartner Advisory, Dataquest (August 7, 2000).
2. *Third Generation Wireless: The Impact on Global Competition*, ICM Global Intelligence (November 1999).

The Rise of the 3G Empire: Even Rome Wasn't Built in a Day

Brian Modoff
Senior Analyst, Wireless Equipment
Deutsche Banc Alex. Brown

Michael Thelander
Associate Analyst, Wireless Equipment
Deutsche Banc Alex. Brown

Over the last several months, the mobile wireless industry has faced a number of challenges that have caused some to doubt whether or not third-generation (3G) mobile services will ever happen. Carriers spent billions of dollars obtaining 3G spectrum without anticipating that the capital markets would contract, making it more challenging to obtain necessary funds to deploy their 3G infrastructures. The worldwide economy has turned south, which has resulted in consumers limiting discretionary spending (cell phones) and wireless carriers refocusing their resources on strengthening their balance sheets. At the same time, the capital markets have contracted, thus limiting carriers from gaining access to reasonable financing terms. Finally, actual data rates are well below expectations and technical challenges and network interoperability problems persist, causing some well-publicized 3G launches to be delayed.

We still believe 3G will happen. However, instead of just providing market segment forecasts and hyping the possibilities of 3G applications, we are going to take a step back and focus on the underlying technologies that give us 3G services.

As depicted in the following chart, regardless of the network technology, the 3G standards are the same. As you will find after reading this report, the paths carriers choose to meet these standards are substantially different.

In *Figure 1*, we provide the target data rates for 3G and how these higher data rates can make download times more reasonable. Just as you can drive faster on an open highway, the higher data rates will only be achieved when network capacity is not being constrained by other subscribers. We should also point out that the download times are not an International Mobile Telecommunications (IMT)-2000 requirement, but are based upon our calculations that assume the minimum required data rate is available throughout the duration of the download.

The Mobile Network: Anatomy 101 of Cellular Technology

In order to understand—and, hopefully, appreciate—the evolutionary path to 3G, you need to start with the basics of telecommunications networks. For, you see, the path to 3G includes a lot more than creating "gee-whiz" handsets, obtaining frequency spectrum and swapping out some radio equipment. Instead, the path involves software and hardware upgrades throughout virtually the entire telecommunications network.

The Radio Access Network

The radio access network (RAN) is comprised of three basic components: the handset, the base transceiver station (BTS), and the base-station controller (BSC). If you fall into the 42 percent of the U.S. population (our apologies to the international readers) who own a cellular phone, then you are already intimately familiar, and in some cases frustrated, with the first component in the radio network. Cell phones are designed for a specific carrier and the frequency band (800 MHz, 1900 MHz) and mode (Global System for Mobile Communications [GSM], time division multiplexing [TDMA], etc.) they use.

Some operators use frequency bands in their networks that vary by markets due to spectrum limitations that prevent them from owning enough spectrum in the same band. If your carrier falls into this category, you need a multi-band (800 MHz and 1900 MHz) phone. A multi-mode phone is also required if you leave the comforts of modern civilization (digital networks) to go fishing for walleye in the Boundary Waters of Northern Minnesota, where they only have analog service.

The next component in the radio network is the BTS. BTSs are available in three sizes: pico (shopping malls, convention centers), micro (highways), and macro (coverage up to a 200 km radius). Generally, pico stations operate in indoor environments and are quite small.

Micro stations are usually deployed within the coverage space of a macro cell to provided additional capacity or along major rural highways. The BTS is comprised of the radio equipment, including processing equipment, filters, cables, antennas, and power amplifiers.

Macro cells are popular with operators because of their wide coverage area, which allows operators to cover a large, low-subscriber-density, geographical region with a relatively limited number of cells. Macro cells are also the most

FIGURE 1

3G—What It Means For You
IMT-2000 Requirements

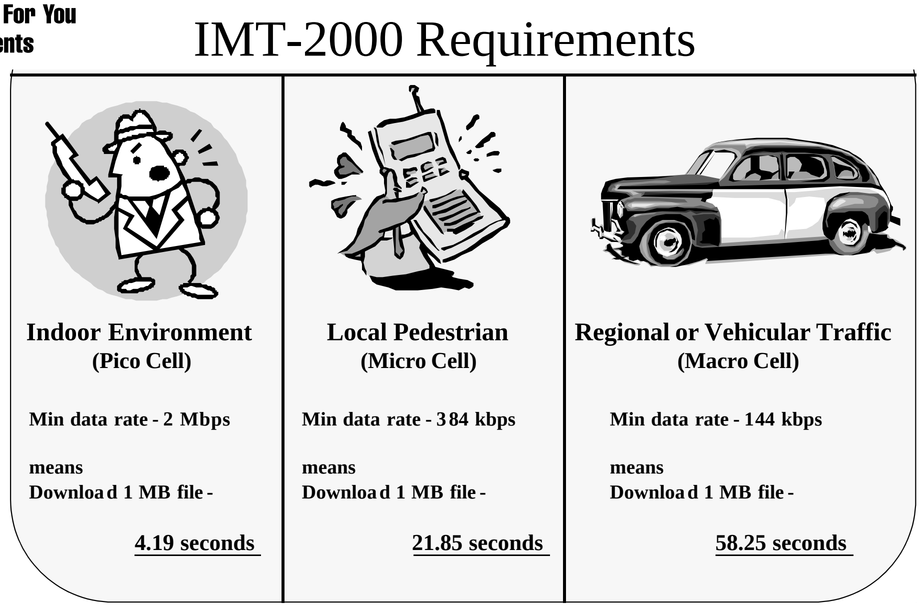

Source: DBAB

visible of the three cell types and can be readily spotted in rural areas—just look for the antenna tower.

The BSC is the final major component of the RAN.[1] In addition to determining the power output between the BTS and mobile phone as well as managing the handover of mobile phones between base stations, the BSC consolidates the transmissions from the BTSs and sends the transmissions on to the mobile switching center (MSC).

The Backhaul

The backhaul portion of a mobile network refers to the transfer of voice and data traffic from the base station to the BSC or between base stations if base stations are daisy-chained together—for example, along a major freeway. There are three physical mediums that are primarily used to transport this traffic: microwave radios, leased T1/E1 lines[2], or fiber-optic cable.

With 2G networks, carriers typically have used one or perhaps multiple T1 or E1 (Europe) leased lines as the physical medium. Even then, the cost could reach 30 percent of network operation costs. However, with the higher-capacity 3G cell sites (up to 10x capacity), the required throughput generated from one cell site can easily reach up to 5 Mbps initially, and we have heard this throughput could increase to as high as 15 Mbps. Daisy-chain a few base stations together, and the data-capacity requirement increases even further. *Figure 2* provides Nokia's view on the future demand for bandwidth in the backhaul network.

Alcatel has also done a number of market studies regarding current and future backhaul requirements. Based upon their studies, European carriers currently require between 1 to 2 Mbps on average with backhauls in denser areas, such as Italy, requiring up to 4 Mbps. One European carrier, who shares their backhaul with another carrier(s), uses more than 2.5 Mbps on average. With 3G, Alcatel's survey of European carriers suggests that initial backhaul demand will be a modest 2 Mbps but could grow to 5–18 Mbps over the next 5–10 years.

The Core Network

The "core network" refers to the portion of the mobile network that is behind the BSC and includes the MSC, as well as the gateways, or entrance points, to other networks.

As mobile networks transition to 3G, changes need to be made to the core network as well, including implementing software upgrades, adding new servers and gateways (connecting points) to other networks, and using new protocols to transport and route traffic through the network. In general, these changes are required to ensure that the core can support applications such as voice, data, and video; higher data rates; and new transport and switching technologies. Wireless operators need to also have the ability to add subscribers while decreasing capital and operational costs.

These networks and their various components provide and/or support a number of functions in the mobile world:

- Billing services
- Phone registration into a home or foreign network
- Call set-up and tear-down
- Interface to the "fixed Internet world" and public switch telephone network (PSTN)
- Real-time voice and data service

Just as importantly, a lot of the debate between code division multiple access [CDMA]2000 and wideband code division multiple access [WCDMA][3] centers around the required changes to the respective core networks.

In the past, the interface between the RAN and the core network was unique to each vendor. As a result, operators were virtually forced to select the same vendor to provide the

FIGURE 2

Mbps for Macrocell Base Station

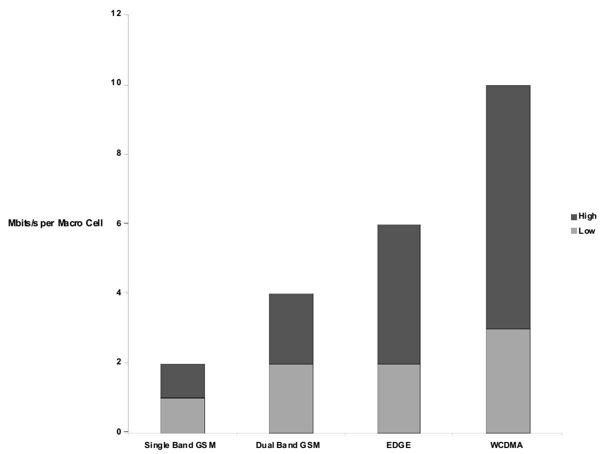

Source: Nokia

RAN and core solution. Now, vendors have accepted an open standard, called IOS v.2 and IOS v.4, which gives operators more options. The third-generation partnership project (3GPP) is trying to take the concept of an open standard to the next level. Their ultimate goal is to have interoperability throughout the entire network—for example, between Vendor A's RNC and Vendor B's Node B.

At the time of publication, original equipment manufacturers (OEMs) and carriers were conducting interoperability tests between the mobile stations and the base stations, with plans to test between the RAN and core network in fall 2001. The 3GPP also intends to have interoperability between 2G and 3G solutions, for example Vendor A's WCDMA solution deployed on top of Vendor B's GSM network. This scenario would support seamless handovers between WCDMA and GSM networks and would also allow more open competition.

All Roads Lead to 3G

The road(s) to 3G vary by technology with different starting points, ending points, and forked roads in between. *Figure 3* shows how the various second-generation (2G) technologies evolve to 3G.

Global System for Mobile Communications

In the beginning (circa 1980s), man created five incompatible analog cellular systems in Europe, and it was good. However, incompatibility problems resulted in a virtual Tower of Babel scenario when cellular subscribers ventured outside of their native countries (networks). In the north, there was NMT 450 and 900. The United Kingdom had TACS and E-TACS, France had Radiocom 2000, Italy used RTMI/RTMS, and TCS as well, and West Germany (yes, there were once two Germanys) had something called C-Netz and D-Netz.

Further, over time, it was realized that these first-generation analog systems would quickly have too many subscribers for the networks, despite the disparity of the different systems. Thus, in keeping with the theme of a unified Europe and to limit research and development costs, a unified *digital* cellular system was needed. As a result, the Conference of European

FIGURE 3

All Roads Lead to 3G

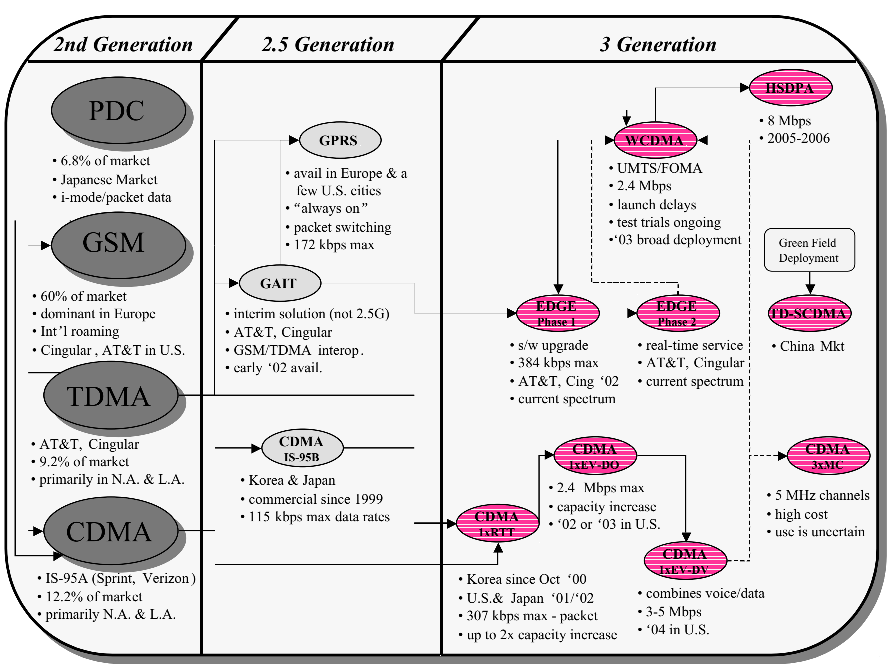

Source: DBAB

FIGURE 4

GSM Network Diagram

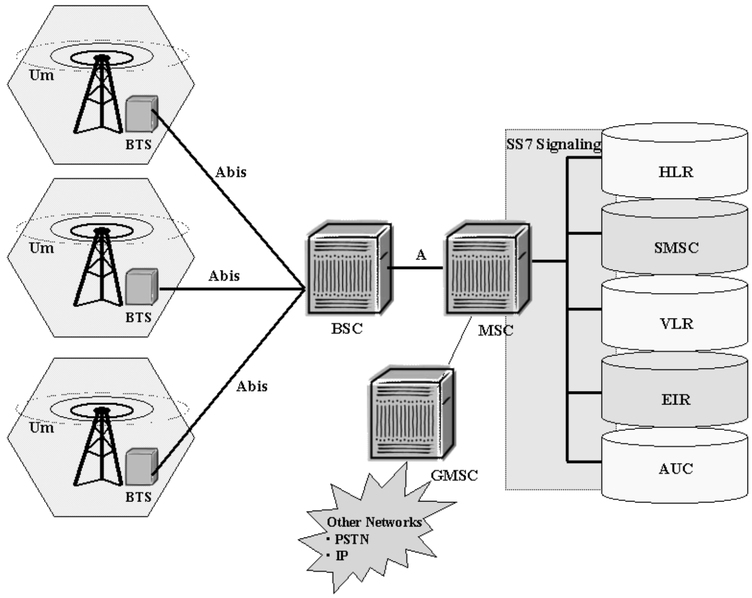

Source: DBAB

Posts and Telecommunications (CEPT) established what became known as GSM,[4] with an objective to develop universal specifications for mobile communications in Europe.

One important aspect of GSM is short message service (SMS), although SMS is also available in TDMA (receive only) and CDMA. SMS is a short message delivery technology that is comprised of a fixed length of characters with the newly defined, but not necessarily deployed, standard supporting a far greater number. Since SMS messages are sent using a separate network (the signaling network and signaling channels on the radio interface), a user can send and receive a SMS while engaged in a telephone conversation.

Currently, nearly 20 billion SMS messages are sent per month, even though it has not really taken off in the United States due in large part to the interoperability between carrier's networks, relatively few SMS–enabled phones, and the lack of customer awareness. However, several carriers, including Verizon (CDMA) and Cingular (GSM and TDMA), are beginning to actively market this feature, and for good reason.[5] In Europe, Vodafone (GSM) reported that nearly 10 percent of its FY00 revenues came from SMS traffic. In addition, because SMS messages are so economical to deliver, SMS is one of the most profitable services offered by wireless carriers.

SMS also has its own evolutionary path that coincides somewhat with the path to 3G. Enhanced messaging service (EMS) will allow new font styles (e.g., bold), simple melodies, pictures, and sounds to be attached to messages. Multimedia messaging service (MMS) follows EMS and will allow short video or audio clips (e.g., MP3, MPEG 4) to be attached. Carriers will likely charge extra for these services.

In *Figure 4*, depicting a GSM network, you can see we have listed the U_m, Abis and A interfaces that define the standards between the mobile station and the BTS (U_m), between the BTS and the BSC (Abis), and between the BSC and the MSC (A). The technical description for one of these interfaces is a standalone document that contains a lot of technical jargon and minutiae that we believe are beyond the interests of any investor. However, after a cursory review of the documents, it should become readily apparent that changing and interpreting (for interoperability) these standards can become a cumbersome task. Here is what we believe you need to know about these three interfaces, starting with the U_m.

***U_m*:** The GSM air interface standard is referred to as the U_m. The GSM frequency spectrum is 25 MHz wide and is divided into 124 individual carrier frequencies along with a 200 KHz guard band. One or more carrier frequencies are assigned to each BTS, depending on how many sectors the cell site contains. The division of frequencies into multiple frequencies, or traffic lanes, is called frequency division multiple access (FDMA). Each of these carrier frequencies subsequently uses a radio network access technology called time division multiple access (TDMA).

TDMA is a digital multiplexing technique that facilitates multiple users accessing a single radio frequency (RF) channel. Imagine a traffic cop merging four lanes of traffic into a single lane of traffic. The cop sequentially signals to the driver in each of the four lanes. When signaled, the driver proceeds ahead without fear of sideswiping a car in an adjacent lane (see *Figure 5*).

With TDMA, a time frame is divided into time slots, in this case eight slots that last for 0.577 milliseconds (ms). Therefore, the time frame is 4.615 ms (8 * 0.577 = 4.615 ms) in length. These slots are allocated to users on the network in order to let the user to "talk" without experiencing interference. Another way of looking at it is that each time slot represents a logical channel that is assigned to a subscriber.

FIGURE 5

Our TDMA Analogy

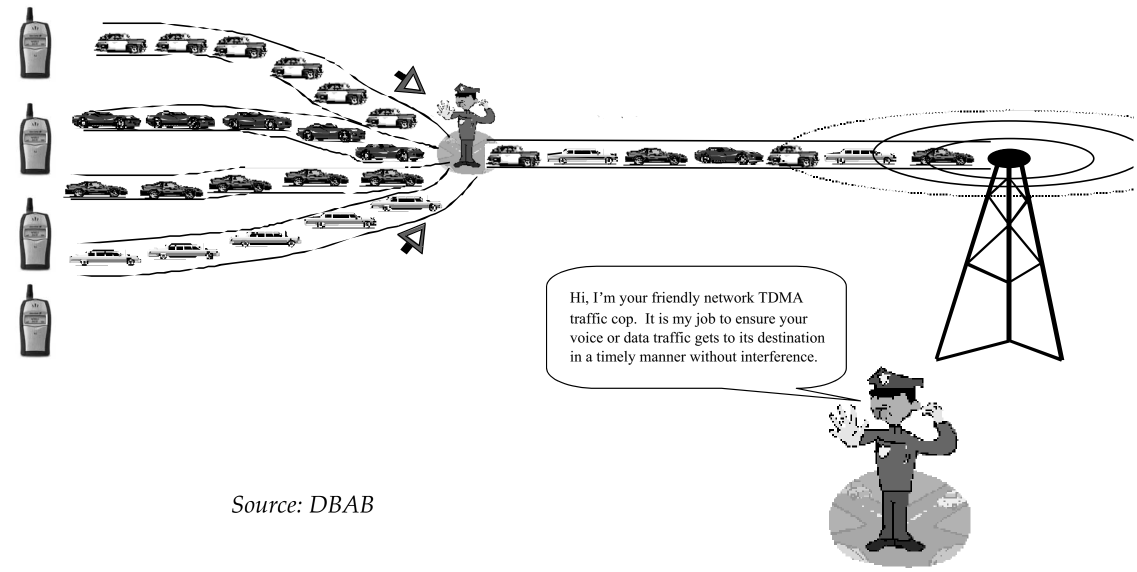

Source: DBAB

The signaling channels, part of the "extended" signaling network, are broken down into three groups. The broadcast channels (BCH) carry downlink traffic only and provide time synchronization and frequency correction information that the mobile device uses to maintain an optimal link with the base station. Additionally, the BCH carries point-to-multipoint SMS messages (one base station to many mobile subscribers). The common control channels (CCCH) handle call notification (downlink) and network access requests (uplink). Finally, the dedicated control channels (DCCH) handle roaming, handovers to other cell sites, and encryption.

While it may seem a bit technical and burdensome to discuss an aspect of a network that you cannot even see, let alone reach out and touch, there is a method behind this madness. First, it should now be apparent why a subscriber can receive an SMS message when talking on the phone—the SMS message is transmitted on a separate channel. We also provided this discussion to illustrate that there is more than basic voice traffic using up the precious spectrum and that the network uses these signals to control each subscriber as they move throughout the network.

A_{bis}: The A_{bis} defines the interface between the BTS and the BCS. Currently, leased T1/E1 lines predominantly provide the physical link for this interface. However, we believe wireless and fiber connections will play a more dominant role with 3G due to the required data link between the two components. Like the U_m, the A_{bis} consists of traffic and signaling channels, with each transceiver requiring its own signaling channel. Since the location of the traffic and signaling channels is not necessarily consistent across various vendors' platforms, incompatibility issues can develop if a carrier mixes and matches its BSC/BTS solution. Of course, this problem is supposed to go away with the 3GPP standard. Finally, the A_{bis} carries operating and maintenance information to and from the BTS.

A: The A interface lies between the BSC and the MSC, which we will discuss in the next paragraph. The A interface uses 32 time slots with a maximum of 30 time slots for traffic and the remaining time slots dedicated for sending control and signaling information. Like the other two interfaces, there are also a number of defined protocols that exist on lower levels of the open system interconnection (OSI) model.

The next major component in the GSM network (or, for that matter, all mobile networks) is the MSC. A MSC is a switch that sets up connections to other MSCs or BSCs. If the MSC also interfaces with another network, such as the PSTN, it is referred to as a gateway MSC (GMSC). A number of interfaces, each with its own defined protocol, extend off the MSC to other MSCs, the signaling system 7 (SS7) network, other networks, or database servers. Each of these interfaces uses SS7–based protocols. We should also point out that some vendors provide an integrated BSC/MSC platform even though our network diagram implies a physical separation of the two components.

The home location register (HLR) is a database that contains user-specific subscriber information. When you first subscribe to a carrier's service, you are registering your phone with the HLR. The HLR also supports charging and accounting features. If you venture into another region of the network (outside of the domain of the HLR) or into another carrier's network, the visitor location register (VLR) is responsible for registering the phone into the network and informing the subscriber's HLR of where the subscriber is currently located. Thus, when a call is placed to the mobile subscriber, the HLR knows where the subscriber is located (which MSC) and routes the call accordingly. Of course, all of this signaling traffic is done using the SS7 network. There is one VLR per MSC, and generally the VLR is integrated with the MSC.

The authentication center (AUC) database generates and stores security-related data that is used in order to prevent fraudulent use by an unauthorized person. The final major database we are mentioning is the equipment identity register (EIR). The EIR contains a list of all valid devices by recording each terminal's international mobile equipment identity (IMEI). If a mobile phone were reported stolen, the EIR would prevent the device from gaining access to the network.

In actuality, a GSM handset is technically subscriberless. Instead, a GSM subscriber is assigned a subscriber identity module (SIM) card that has its own unique international mobile subscriber identity (IMSI). The SIM, or smart card, contains all of the user's personal information, such as security information, address book, and assigned telephone number. Thus, any SIM–enabled phone becomes a personal phone once the SIM card is inserted. Although currently available in Europe, SIM technology is making its way to the United States and is being incorporated into CDMA and TDMA protocols.

SIM cards are also a cheap way to provide worldwide global roaming. Instead of buying a tri-mode GSM phone, subscribers can just take their SIM card to Europe and rent a phone. With the SIM card installed, the subscriber can receive calls just as if he or she had their personal phone. Callers would also be oblivious to the fact that the SIM cardholder was in another country.

General Packet Radio Service

General packet radio service (GPRS) is the next step along the GSM evolutionary path to 3G and is often referred to as 2.5G.

Perhaps the most talked about (and overhyped) advantage of GPRS is increased data rates over 2G solutions. GSM networks are limited to 9.6 kbps for data, since data and voice are digitized using the same technology and subsequently transmitted over the same path (circuit). While 9.6 kbps is more than adequate for even the fastest-talking research analyst, it pales in comparison to data rates that we have grown accustomed to while at work or at home (even if you lack broadband access). GPRS, on the other hand, features speeds of up to roughly 115 kbps (theoretical maximum of 172.2 kbps), although actual speeds will be substantially less, on the order of 20 to 40 kbps.

Recall that in GSM radio access networks, each mobile subscriber is assigned one channel, or time slot, in which to transmit data in the upstream and the downstream directions (TDMA). Further, this time slot is allocated to the user even if data is not being transmitted—hardly efficient use of spectrum.

In GPRS, a mobile subscriber can be assigned multiple time slots for data traffic (still using TDMA) in each direction. Initial deployments will support several time slots in the downlink and only one time slot in the uplink (Note: Web browsing requires greater throughput in the downlink). If the channel (time slot) is being used for data traffic, the channel is called a packet data channel (PDCH). The ratio of PDCH to "regular" voice channels is based upon current demand for both channels. Further, the data packets are interleaved among the voice traffic for further efficiency.

There are also a number of packet-signaling channels. In summary, these signals provide point-to-point and point-to-multipoint (downstream) network information, such as network timing and notification of an arriving packet.

Our final comment on the GPRS RAN concerns data coding. There are four data coding/modulation schemes that are used, depending upon network conditions. Under ideal conditions, the fastest technique, called convolutional coding, can achieve a data rate of 21.4 kbps per time slot. If a carrier is generous enough to allocate all eight TDMA time slots to a subscriber, one could obtain the theoretical maximum GPRS data rate of 171.2 kbps (8 * 21.4 = 171.2). However, just as driving on the Autobahn does not guarantee that you can play Mario Andretti, it should also be obvious why you will never achieve the maximum GPRS data rate. Instead, average data rates will probably be between 20 and 40 kbps.

Early GPRS handsets also have a throttle on them that limits the number of time slots they can support. Although a number of handset OEMs claim their GPRS handsets can support 115 kbps, the published specifications imply otherwise by stating that the handset supports four slots in the downlink and one slot in the uplink. A few handsets only support two slots in the downlink.

Figure 6 shows how data rates (Y axis) increase as the number of assigned time slots (X axis) increases and the data-coding scheme (Z axis) is changed. The upper right corner of the chart represents the theoretical maximum data rate, while the white circle represents reality—20 to 40 kbps.

At first glance, actual GPRS data rates may not sound like a lot of bandwidth, especially when compared to the maximum data rates that are often quoted in the press. However, keep in mind that nearly 80 precent of U.S. households that use the Internet access the Internet using a 56 kbps modem, and even then actual data rates are almost always lower. Thus, GPRS can potentially deliver practically the same bandwidth to the mobile device that most U.S. households have grown accustomed to at home. Further, Web applications on a mobile handset require less bandwidth than their fixed wireline counterpart. As a result, actual GPRS data rates should suffice for some applications, or at least be an improvement over 2G networks.

We believe an even more compelling reason than the speed advantage is that GPRS is "always on," meaning that your phone is always connected to the network and ready to receive and notify you of incoming messages (trade confirmations or stock-price alerts). This is also called immediacy. Further, carriers can now charge you based upon how much data you send or receive, versus on a per minute basis. For example, Japan's NTT DoCoMo charges $0.0024 for 128 bytes of data. In Europe, Vodafone was recently charging $5.66 a month plus $0.03 per kilobyte of data downloaded. AT&T Wireless is charging $50 per month for 1 MB and 400 voice minutes per month, with additional data packets priced at less than $0.01 per kilobyte. Finally, Cingular, which launched service in Seattle in late August 2001 and plans to have coverage in all of its GSM networks in 2002, offers an entry-level data package for $14.99 that includes 0.5 MB of data and 100 messages.

Although your phone is always connected to the network, it is not tying up a radio channel unless data is being sent

FIGURE 6

Varying Parameters and Sample GPRS Data Rates

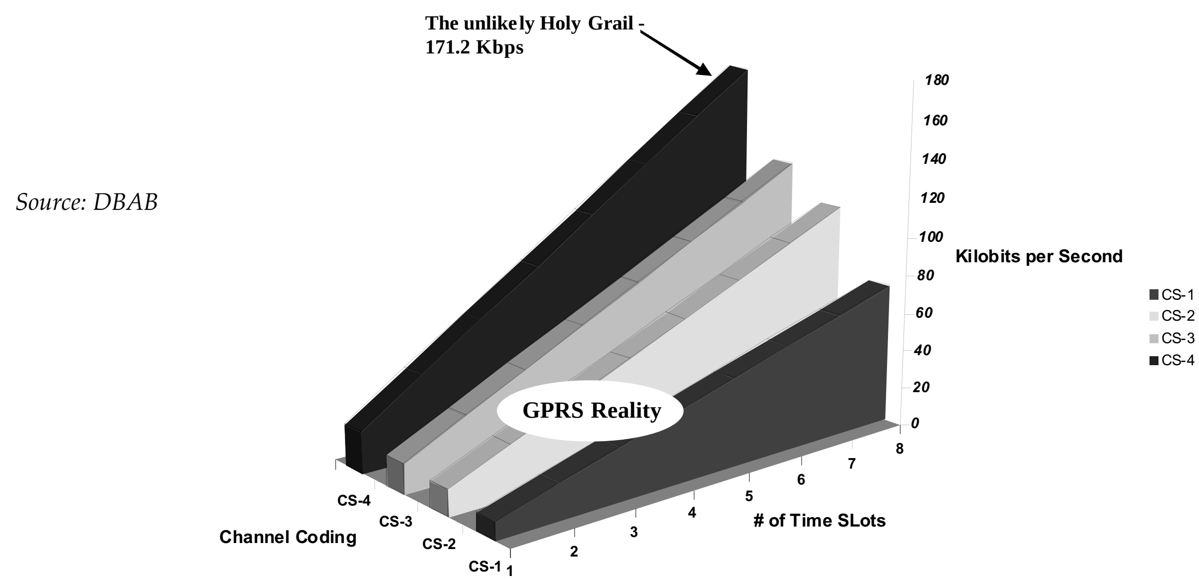

to or from the GPRS handset (due to packet versus circuit switching). Thus, network capacity is more efficiently used (for data traffic only) with GPRS technology. GPRS does not address voice capacity in mobile networks. However, GPRS can be a heavy user of GSM voice capacity if not properly controlled.

Two additional nodes are added when GPRS is overlayed on a GSM network. The serving GPRS support node (SGSN) is responsible for routing incoming and outgoing Internet protocol (IP) packets that are from, or intended for, any GPRS subscriber within the area served by the SGSN.

In addition, the SGSN provides encryption/decryption, session management (call set-up), mobility management (cell handover), and connection to other nodes. If a mobile subscriber moves to a location in the network that is supported by a different SGSN, the new SGSN requests information from the old SGSN and, if necessary, notifies the HLR, VLR, and MSC (of course, using the SS7 signaling network).

A gateway GPRS support node (GGSN) is the interface between the GPRS network (see *Figure 7*) and other packet data networks. The GGSN converts packets received from the SGSN into the appropriate protocol (e.g., encapsulates a foreign packet with an IP protocol for IP networks) that is required by the network receiving the GPRS packet. Conversely, the GGSN receives packets from other networks, coverts them to GPRS format, and sends the packet to its intended SGSN. In general, one GGSN connects to multiple networks and has connectivity with multiple SGSNs in the GPRS backbone network. The GGSN also assigns a packet data protocol (PDP) address when requested by the mobile subscriber (requested through the SGSN). The PDP includes the newly assigned IP address as well as the quality of service (QoS) and address of the GGSN that is connected to the targeted packet data network.

The actual number of deployed SSGNs and GGSNs is dependent upon the number of active data subscribers and the capacity limits of the vendor's solution.

Enhanced Data Rates for GSM and TDMA Evolution (EDGE)

For some carriers, EDGE represents the next step in the GSM/TDMA evolutionary path to 3G. In fact, by definition, EDGE is a 3G service that could become the final resting-ground for some carriers—at least until 4G (whatever that is) comes around. EDGE is broken down into three areas: standardization of the physical layer, protocol changes for enhanced circuit-switched data (ECSD) and enhanced GPRS (EGPRS), and two phases. EDGE Phase 1 is concluded with Release '99.[6]

As you probably noticed, the transition from TDMA/GSM to GPRS was not a simple task, and a number of new components were added to the network. Fortunately, the transition to the first phase of EDGE is more like a little schoolyard skip than a giant leap.

In fact, the only part of the network that is affected in Phase I is the radio network, due to the introduction of a new, higher-order modulation technique called eight phase shift key (8PSK). The rest of the network, assuming it made the transition to GPRS, already has enough capacity and processing power in reserve to handle the increased data rates.

In order to have higher bit rates per time slot without impacting the underlying channel structure (maintain GSM/GPRS service), the modulation type needs to be changed. EDGE uses 8PSK. 8PSK is a linear modulation scheme in which three consecutive bits are mapped onto one symbol in the I/Q plane. This rate is divided over eight time slots and, with interleaving, ciphering, and burst formatting,

FIGURE 7

GPRS Network Diagram

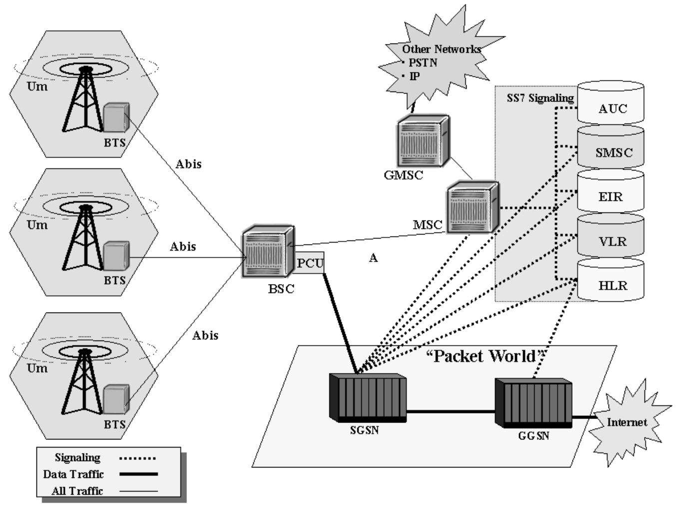

Source: DBAB

it results in a radio data rate of 59.2 kbps per time slot. 8PSK varies the waveform's phase (versus the amplitude or frequency) to one of eight positions. Thus, each symbol can support three data bits, and its theoretical bandwidth efficiency is 3 bits/sec/Hz versus 2 bits/sec/Hz. Under ideal conditions, the fastest channel-coding scheme is 59.2 kbps. If you are allocated all eight time slots, you can achieve the maximum EDGE rate of 473.6 kbps (8 x 59.2 = 473.6).

Figure 8 illustrates how EDGE data rates increase (Y axis) as more time slots are assigned (X axis) and a higher modulation and coding scheme (MCS) is used (Z axis). The first four coding schemes (MCS 1–4) use GMSK, and the remaining five schemes (MCS 5–9) use 8PSK modulation. As in the GPRS chart, the upper right corner represents the theoretical maximum data rate if eight time slots are allocated to one subscriber. If we assume that no more than four time slots are assigned, the typical user data rates should be between 36.2 and 236.8 kbps. Notice that we have included all nine coding schemes, since EDGE improvements to the physical layer of the RAN should support more widespread use of higher order modulation schemes than with GPRS.

The EDGE standard defines link quality control (LQC) procedures that are used to select the optimal channel-coding scheme based upon the quality of the radio link in order to provide the maximum data rate. In contrast to GPRS, EDGE allows the retransmission of previously sent data blocks using a *different* channel-coding scheme than the original scheme. Additionally, EDGE can determine the quality of the radio link more quickly than with GPRS. As a result, a higher-modulation scheme (e.g., MCS-9) can be used, regardless of the radio link quality. Unreadable packets are then retransmitted using a different coding scheme until decoding is successful. According to Ericsson, this approach means that for a given data rate per time slot, EDGE has better coverage than GPRS. This approach is comparable to always talking very quickly but having to repeat oneself on occasion until the listener understands what is being said.

One of the challenges with EDGE technology is being able to incorporate multiple modulation schemes (GMSK and 8PSK) into the same handset and still provide a relatively low-cost, compact phone that has acceptable battery life. The handset can also heat up and battery life can also be diminished if the handset transmits for an extended period of time (just like a machine gun barrel can overheat), as is often the case if more than one time slot is allocated to the subscriber. This is true in GPRS as well. The performance and reliability of overheated components are also impacted (recall that some early GPRS handsets overheated). One major handset OEM even informed us that it will be several years before an EDGE handset will be able to support more than five time slots due to heat dissipation issues.

Although EDGE Phase I provides high data rates, Phase I does not support real-time services, such as videoconferencing and voice over IP (VoIP), nor does it provide the same QoS classes as the Universal Mobile Telecommunications System (UMTS). EDGE Phase 1 is concluded with Release '99.

3GPP is currently addressing EDGE Phase II, which was previously under the purview of European Telecommunications Standards Institute (ETSI). The scope was also expanded to include the alignment with UMTS and the provisioning of IP multimedia services. As a result, Phase II uses an interface to an all–IP 3G core network that is common with the UMTS core network. We will discuss the UMTS core network in a following section along with providing a network diagram that shows the primary components and interfaces.

The Phase II RAN, called GERAN (GSM EDGE RAN) supports several bearer classes (service levels), including maximum bit rate (analogy—being able to drive as fast as 100 MPH when possible) and guaranteed bit rate (analogy—always being able to drive 50 MPH). Header compression (less overhead) is used to achieve higher performance, QoS, and more efficient multiplexing. As a result, more than eight users can share the same carrier frequency. More importantly, half-rate speech channels, in which the voice packets

FIGURE 8

Varying Parameters and Sample EDGE Data Rates

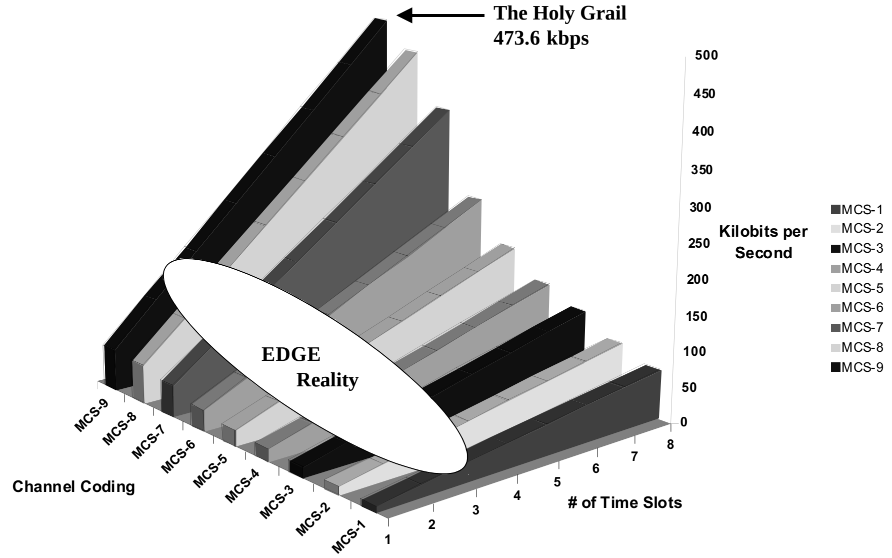

Source: DBAB

are compressed, can be used to achieve higher voice capacity. Finally, GERAN also allows for faster handovers between cell sites.

A Word About Vocoders

In order to process a voice signal for transmission, the analog speech needs to first be converted into a numeric representation of the voice signal. This numeric representation is then processed by the digital circuitry of the handset before it is ultimately converted back to analog and transmitted over the airwaves. A vocoder is the hardware/software solution that is used to convert analog speech into a digital bit stream. There are a number of coding schemes that can be used. However, the 3GPP (WCDMA) has approved the use of adaptive multi-rate (AMR), which is a coding scheme that compresses the voice at varying bit rates—lower bit rates increase the number of concurrent users. The 3GPP2 (CDMA2000) is using what they call selectable mode vocoder (SMV). Like the AMR solution, the SMV can operate in several modes depending on the quality of the network and the desired voice quality. The SMV can increase the capacity of the network (up to 75 percent) if a lower coding rate is used. When AMR is available next year, Nokia believes it will increase GSM capacity by approximately 125 percent. For either solution, voice quality could be impacted—much like recording a television show at a compressed rate (longer tape life) may impact the quality of the recording.

Code Division Multiple Access in a Nutshell

Although it may be easier to say that CDMA is "magic" and leave it at that, we thought we could do one step better without providing Dr. Irwin Jacob's (Chairman and CEO of Qualcomm) private telephone number. Although CDMA and WCDMA are perceived to be entirely different solutions, the truth of the matter is that at a high level, there are more similarities than differences, with Qualcomm owning most of the intellectual property.

CDMA assigns a distinct code to each subscriber. Probably the best analogy is to imagine a room full of people from all over the world (say Finland, Sweden, and San Diego), each having a conversation with a colleague in their respective language. Although the conversations are occurring simultaneously and within the same frequency spectrum, the listener is able to filter out the unintelligible languages (noise) and only hear and understand the conversation that is meant for him or her (see *Figure 9*).

This code is used to spread the information signal. Since the bandwidth of the coded signal is much larger than the original signal, the term spread spectrum modulation is often used when talking about CDMA. Each code is also orthogonal (no correlation) to other codes, thus multiple coded signals can be transmitted over the same time and frequency domain.

Generally speaking, and we emphasize the word "general," WCDMA and CDMA have some common characteristics. In particular, both use a coding scheme to encrypt the data, and both use control channels to manage the network (signaling, power output, network access request) and traffic channels that are used for the data. However, WCDMA and CDMA2000 are not compatible. In addition to using 5 MHz channels, the WCDMA chip rate is different than its CDMA2000 brethren. WCDMA uses a chip rate (how fast the data is spread) of 3.84 Mchips/s in comparison with CDMA2000 that uses a chip rate of 1.2288 Mchips/s for a 1.25 MHz channel. Although it would appear that the higher chip rate and wider channels (5 MHz versus 1.25 MHz) would indicate higher data throughput using WCDMA, this is not necessarily the case. Instead, characteristics such as the channel structure, power control, and synchronization controls play a greater role.

WCDMA networks are asynchronous, which means the base stations do not have to be exactly in time sync in order to hand off—called a handover—a subscriber from one base

FIGURE 9

Our CDMA Analogy

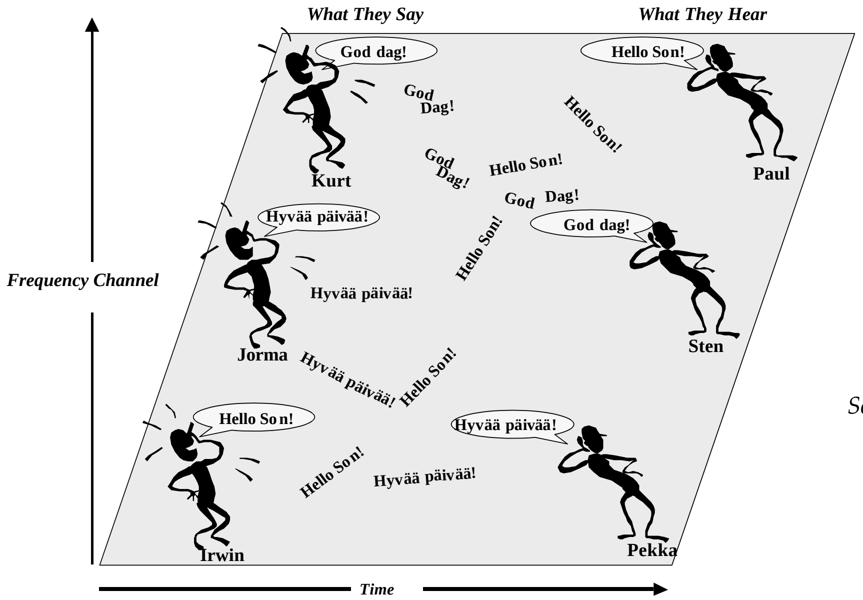

Source: DBAB

station to another (CDMA2000 BTSs are time synched using GPS). Instead, the mobile devices measure the timing differences (sent on a control channel) between two base stations. This difference is reported back to the active base station that can then make an adjustment before switching the mobile device to the new base station.

As a result of the asynchronous and synchronous differences, the WCDMA and CDMA2000 codes are generated and assigned differently. Specifically, WCDMA base stations are separated by 512 unique scrambling codes, while CDMA2000 base stations are separated by a timing offset of the same sequence. Another difference is that WCDMA frames, used to "group" the voice/data/signaling traffic in the RAN, are of length 10 msec and CDMA2000 frames are of length 5, 10, and 20 msec.

Handovers occur between cell sites, frequencies, and across technologies. For example, the network may assign to a mobile phone a new frequency in hot spot areas where there are multiple carrier frequencies being used and interference develops. Additionally, if the phone moves outside of an UMTS network, the phone needs to switch to GSM or GPRS mode. Thus, the multi-mode phone needs to operate in UMTS mode while simultaneously capturing the synchronization information that is being transmitted in GSM mode. A multi-mode phone actually benefits carriers because it allows them to initially deploy WCDMA in a limited area and still provide universal coverage using their GSM network. Nokia has publicly commented that it should have a multi-mode (UMTS/GSM) handset in the second half of 2002, although we remain skeptical. In Japan, a single-mode phone will be offered, since WCDMA will be more widely available.

One interesting feature of CDMA (WCDMA and CDMA2000) is the concept of soft handovers (see *Figure 10*). Unlike GSM or TDMA, adjacent cell sites in a CDMA network use the same frequency. Typically, a mobile device is instructed to switch cell sites and frequencies when the signal strength from the adjacent cell exceeds the strength of the currently assigned cell. The switching between different frequencies is called a hard handover and can also occur within the same cell site if a new frequency is required—for example, if interference develops when using original frequency. However, in a CDMA network, the same frequency can be maintained when switching between cell sites.

While soft handovers are clearly advantageous in a network, it does introduce overhead. Specifically, a handset that is communicating with two base stations can limit the number of additional subscribers that can access the network. Although the amount of overhead can vary, we have come across percentages that range between 25 percent and 35 percent or even higher, regardless of it being a WCDMA or CDMA2000 network.

When the CDMA phone is in a soft handover state, the phone is simultaneously transmitting and receiving from more than one cell site. The BSC keeps track of the signal from the mobile coming from multiple BTSs and selects the best quality signal. As a result, if a packet from one BTS is "clobbered," the same packet is picked up via the path from the other BTS. Thus, the near-instantaneous break in transmission associated with a hard handover is reduced to zero.

In receive mode, the CDMA handset actually combines the multiple signals, each with a slightly different amplitude and time shift due to the different distances traveled by multiple signals, into one cohesive signal.[7] A comparable example would be a disk jockey who uses two turntables to fade out/in music from different sources in order to continuously keep the music playing.

FIGURE 10

Soft Handover in a CDMA Network

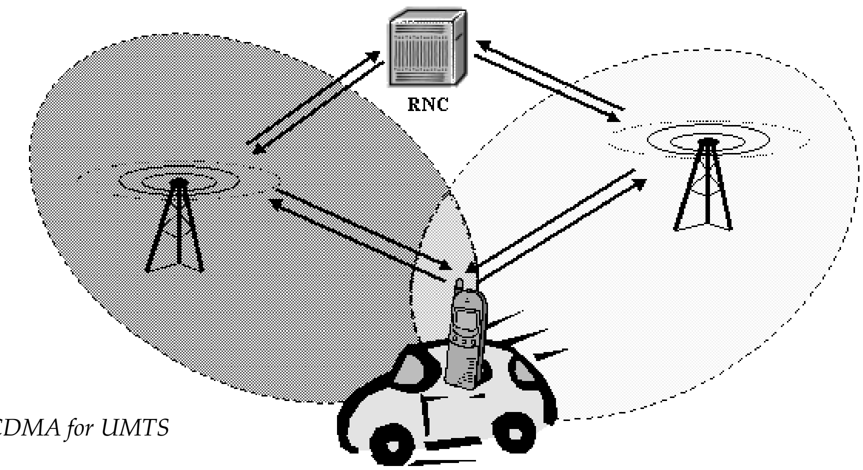

Recreated from WCDMA for UMTS

WCDMA (UMTS) Networks

The Holy Grail for GSM–based carriers, at least those not considering CDMA2000, is WCDMA. To provide UMTS service, the RAN and the core network require changes in order to provide much higher data rates, enhanced QoS, and real-time service. In UMTS, the radio access network is called UTRAN (UMTS Terrestrial Radio Access Network).

The UTRAN includes a frequency division duplexing (FDD) mode and a time division duplexing (TDD) mode. The FDD mode is based on pure WCDMA, while the TDD mode includes an additional TDMA component. The air interface standard used by UMTS is called U_u. The U_u manages connections and handles signaling, paging, and timing information. At the same time, the U_m (GSM) interface remains in order to provide backward compatibility for the GSM phones and to provide coverage in areas where UMTS has not been deployed.

WCDMA uses 5 MHz channels, or multiples of 5 MHz. However, the current spectrum used by GSM in some markets lacks the room to provide UMTS service, since most of the spectrum has already been licensed to provide other services. Additionally, a guard band is required to prevent out-of-spectrum emissions. As a result, UMTS typically operates in newly licensed spectrum (prior to publication, the U.S. spectrum had not been defined). Even then, spectrum is a limited resource, so it did not make sense to define a standard that required more than 5 MHz. It was also felt that 5 MHz was adequate to provide the data rates required to provide 3G services

Chip Rates, Codes, Data Rates, and Coverage

In our discussion on GPRS and EDGE, we suggested that theoretical maximum data rates might never be achieved due to lower modulation and coding schemes being used and not all eight time slots being assigned. In the WCDMA (and CDMA2000) world, maximum user data rates (2 Mbps) are possible, but not probable. The reasons, however, are entirely different, and a bit more complicated.[8]

In CDMA networks, the limiting factor is the processing gain. Processing gain is expressed as a function of the spreading chip rate (3.84 Mcps or 1.2288 Mcps) and the user bit rate (for example, 12.2 kbps for voice). In CDMA networks, it is even possible to detect and decode a signal that is below the noise level, which is akin to being able to hear a whisper in a room of loud-talking people. However, as the user bit rate increases for a given chip-spreading rate (e.g., 3.84 Mcps), the processing gain is lowered.

The maximum user data rate for one assigned code is between 400 and 500 kbps. In order to increase the data rate further, additional codes are assigned, effectively creating a parallel data channel to the user. The maximum number of codes is six, which, even after data retransmission, can easily provide a user data rate of 2 Mbps (6*400 > 2000). In CDMA2000 networks, the lower chip rate requires additional codes and a higher modulation scheme to achieve the maximum data rate. WCDMA also uses a higher order modulation in the downlink (quadrature PSK—think four) than in the uplink (binary PSK—think two) to increase throughput in the downlink for a given number of codes. While additional codes can increase the data rate, it also increases the data rate, which requires additional processing requirements (e.g., decoding) for the handset and can reduce the battery life.

We believe operators will initially use lower kbps voice traffic as the determining factor for a WCDMA cell site's area of coverage. Operators will also consider the uplink, since mobile devices do not transmit at the same power level as base stations (on the order of 0.125 watts maximum versus 10–40 watts).

A number of factors can influence cell-site coverage and capacity, and we want to stress that the numbers we cite may not reflect actual values in a deployed network. However, on a relative basis, we believe that the numbers provide good insight into 3G capacity versus coverage.

Nokia calculated, based upon a number of assumptions, such as tower height and carrier frequency, that the cell range (for this scenario) in an urban environment for WCDMA voice traffic (12.2 kbps) is 2.3 km. Using the same assumptions, the maximum range for 144 kbps (indoors) is 1.4 km. We point out that AMR speech codec can increase the range for voice traffic with the trade-off being reduced voice quality.

Under another scenario, Nokia calculated the transmission range in the uplink for different data rates. As *Figure 11* illustrates, the uplink range at 2 Mbps is approximately one-third the range at 32 kbps. The reduced range is due to the lower processing gain achieved at higher user data rates. It should be noted that, at least initially, high data rates (>32 kbps) in the uplink may not be necessary and that it will be the downlink data rate that matters the most.

In the downlink, the base station transmits at a much higher power output than the mobile station can transmit in the uplink. As a result, Nokia states that a cell site can provide 1 Mbps rates in the downlink to any user who is within range to place a call (the uplink is the limiting factor). However, this would require high amplifier power and that no other users were utilizing capacity in that cell. These results suggest that 2 Mbps user data rates are possible in a 3G network, especially in a pico or micro cell. However, as Nokia points out, it would require that there are no other subscribers using capacity in that cell.

Additionally, if operators reuse their GSM 1800 MHz cell sites as WCDMA sites, Nokia estimates that 144 kbps data rates are possible throughout the cell. For GSM 800/900 MHz, this user rate drops to 64 kbps. Thus, operators do not necessarily have to increase the density of cell sites if they are content with the lower data rates.

Ironically enough, although maximum data rates decrease as the mobile device is transported more quickly in the network (e.g., in a car), coverage efficiency in a moving vehicle actually increases due to a phenomenon called fast fading.

UTRAN Infrastructure

The UTRAN introduces two new network components: the RNC and Node B.

Node B replaces the BTS. Specific Node B tasks include the following:

- Mobile handset control and communication
- Error correction
- Spreading and despreading of the signal
- Measuring the strength of the signal
- Providing power adjustment information to the mobile device

Node B is generally co-located with a GSM BTS and can support both FDD modes of operation. No vendor to date has announced support for TDD. However, the need for higher data rates (e.g., more than 64 kbps) may require operators to deploy additional cell sites, especially at 800/900 MHz.

We believe that UMTS will initially be deployed in "hot spot" regions and that ubiquitous 3G service will not be available for a number of years beyond the initially limited 3G deployments.

The radio network controller (RNC) is similar to the BSC. Specific tasks include the following:

- Resource management of the radio resources
- Protocol exchanges between its interface with the core network and the UTRAN
- Operation and maintenance of each cell (called a radio network system)
- Handovers between cell sites

What would a new technology be without new interfaces that have catchy and intuitive names? WCDMA introduces the Iub (between RNC and Node B—manages the common channels and controls synchronization and radio links), the Iu (–CS for circuit-switch data and –PS for packet-switch data—

FIGURE 11

Data Rate versus Transmission Range in the Uplink

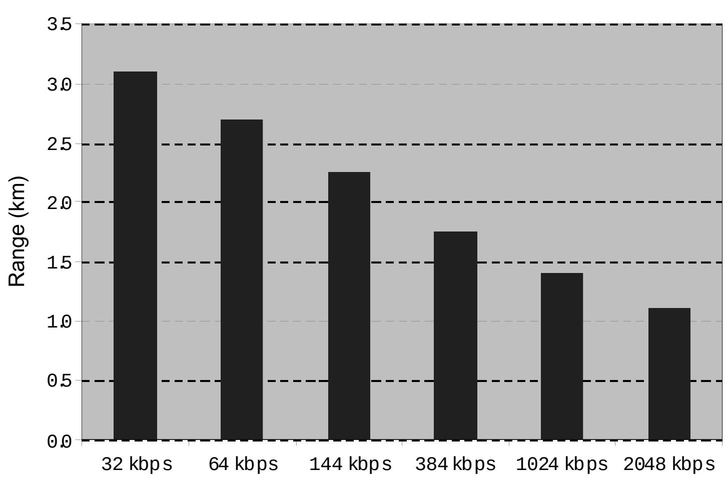

Source: DBAB

note the voice and data are treated differently), and the Iur (RNC to RNC). All three initially use an ATM transport.

The UMTS core network uses an all–IP network with Release 5. In addition to some of the interfaces in *Figure 12* that impact the 3G core, a media gateway (MGW), if not already introduced into the network, is added as well. We should note that a MGW is not a "3G device," since it can be installed during any portion of the upgrade process and is also used in applications outside of mobile networks.

A MGW resides at the boundary between different networks (e.g., PSTN, TDM, ATM, IP) and is used to bridge the gap between transmission technologies by providing protocol conversion, QoS mapping, and voice encoding/decoding.

A MGW provides multiprotocol label switching (MPLS), real-time routing, VoIP, and Gigabit Ethernet termination. A separate media gateway controller (MGC) is used for call control so that the call control "intelligence" is handled outside the media gateways (more room for traffic). Eventually, the media gateways will completely replace functionality of the MSCs when an all–IP core is finally realized.

The home subscriber server (HSS) is an expanded HLR that maintains user profiles, user tracking, and services to which the user has subscribed. The HSS also provides the mobile IP functions that we discuss in more detail in our CDMA2000 section (no sense in having the same discussion twice).

The final new functionality worth mentioning is the call state control function (CSCF). The CSCF handles responsibilities such as call set-up/termination, address analysis, and address modifications, and it interacts with the HSS to obtain user profile information.

CDMA Networks (IS–95 to CDMA2000)

We could just tell you that the evolutionary path from cdmaOne (IS–95) to 1xRTT requires new channel cards and a software upgrade into the base station. Instead, we thought we would delve a bit deeper into the technical changes that result in the higher performance of the CDMA RAN as it evolves to 3G. Note that the chip rate remains at 1.2288 Mchips/s per 1.25 MHz band. One other interesting note is that all CDMA handsets, regardless of generation, are forward and backward compatible with the next-generation networks. What this means is that a cdmaOne handset works on a CDMA2000 network just as a CDMA2000 handset works on a cdmaOne network. We believe this is a distinct advantage in the evolutionary process.

IS–95 A/B

- IS–95A uses 64 unique codes (called Walsh codes) that encrypt the transmissions on the forward link (from base station to mobile phone). This equates into 55 traffic channels plus a pilot channel, seven paging channels, and a synchronization channel.
- IS–95B, like IS–95A uses a vocoder (voice coder) at 9.6 kbps or 14.4 kbps. At 14.4 kbps, the voice signal has higher quality, but the trade-off is reduced network capacity. Depending upon traffic requirements, an IS–95B base station can assign a mobile unit up to seven supplementary channels, each supporting 9.6 kbps or 14.4 kbps. When combined with the original channel, the theoretical maximum 115.2 kbps is achieved (8x14.4 = 115.2). However, current data rates in Korea, where IS–95B is deployed, only achieve 64 kbps.
- Other IS–95B enhancements include better power-control capabilities and improved handoff features that

FIGURE 12

UMTS Network Diagram during Transition

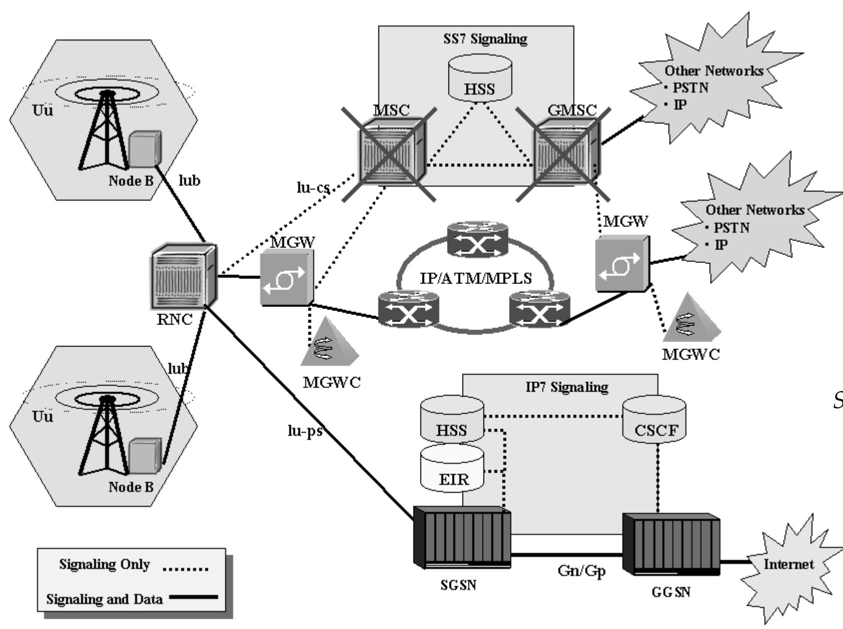

Source: DBAB

are intended to decrease dropped calls and reduce soft handoff times.

1xRTT

- New transceiver cards at the base stations and a software upgrade at the BSCs are required to support higher-speed packet data traffic. An upgraded handset baseband chip is also required.
- 1xRTT uses 128 unique Walsh codes, which theoretically double the capacity of each 1.25 MHz carrier channel. Recall that IS–95 uses 64 codes (64x2 = 128). However, interference is still the limiting factor with CDMA, so doubling of number of codes does not double the radio access network capacity.
- 1xRTT introduces changes to the RAN channel structure. The changes include extending the pilot channel to the reverse link as well increasing capacity in the reverse link, faster power control to improve network capacity and reduce power consumption at the base station, and a faster paging channel that can reduce the power consumption of the mobile device.
- In order to support packet services, 1xRTT can dynamically assign additional supplementary code channels in the forward and reverse links.

1xEV-DO

- 1xEV-DO is a data-only solution that separates data onto its own radio channel. In order to support the higher data rates, additional base stations may be required. 1xEV-DO also requires a new channel card and software upgrade to handle the higher data rates.
- Handsets need to have an upgraded baseband chip.
- 1xEV uses 8PSK instead of quadrature PSK as its modulation scheme. This increase is comparable to the higher modulation scheme used by EDGE versus GPRS in which the maximum data rate triples.
- 1xEV-DO uses a separate 1.25 MHz channel. If voice traffic is required, a 1xRTT or IS–95 channel is used.

Unlike the other CDMA2000 phases that we discussed, the 1xEV-DV (data and voice) standard is still being debated. There are two proposed 1xEV-DV standards. Motorola, Philips Semiconductor, Nokia, and Texas Instruments are backing their 1xTREME proposal. At the same time Lucent Technologies, LG Electronics, LSI, Qualcomm, and Samsung support their competing solution. One of the primary differences between the two proposed solutions is the modulation technique used in the RAN. The 1xTREME camp is proposing using 64 quadrature amplitude modulation (QAM), while the other proposed solution uses 16-QAM and MIMO technology.

With QAM, the amplitude and the phase of the signal are modulated to increase throughput at the sacrifice of being less resistant to noise and a higher probability of dropped packets. If you recall our modulation diagrams, just imagine there are 64 dots (64-QAM) or 16 dots (16-QAM). We note that with 64-QAM there are 64 discrete states or 5 bits ($2^5 = 64$). You may also remember that MIMO technology uses multiple antennas at both locations in order to boost data rates, thus a higher modulation scheme may not be required. We favor MIMO technology over increased modulation (always a tough call—pun intended—in a mobile environment).

CDMA2000-3X

- Three 1.25 MHz channels are grouped together in the forward direction to create a 5 MHz super channel.
- Backward compatible with 1xRTT and IS–95.
- Higher voice capacity than 1xRTT.

Given some of the challenges and high costs associated with CDMA2000-3x, carriers may remain with 1xEV-DV and not deploy CDMA2000-3x. Another reason cited by Verizon Wireless is that 3xRTT does not offer any improvements over 1xEV-DV, yet it requires the same amount of spectrum as WCDMA. Thus, Verizon Wireless would more likely deploy WCDMA than 3xRTT, although that scenario is also unlikely. We do point out that the 1xRTT standard incorporates 3xRTT, meaning the 3x standard has been accepted.

CDMA2000 Core Network

A CDMA2000 network is displayed in the *Figure 13*. Chances are the network looks very similar to the UMTS network. For example, both have an MSC, an HLR, a VLR, and some sort of packet node. In fact, the role of the equipment is virtually identical in both networks. Like they say, "looks can be deceiving."

The critical difference between the two networks lies in the signaling software protocols and how information is exchanged around the network.[9] The CDMA and TDMA core networks are based on the ANSI–41 standard, while GSM uses GSM MAP. However, although the GSM MAP and ANSI–41 protocols are not compatible, the SS7 network in both instances is used to transport the signaling information.

Signaling is paramount to the mobile network. It verifies the authenticity of subscribers and gives them access into the network. Signaling is also used to set up calls, keep track of the location of subscribers, send SMS messages, and the list goes on. The same holds true in an IP environment.

CDMA2000 introduces new hardware in order to provide mobile data service in an IP environment. The packet data server node (PDSN) has functions that are similar to those performed by the SGSN. However, the PDSN is purely IP–based networking, whereas the SGSN also handles signaling with the SS7 world.

PDSN is the "liaison" between the mobile device and the mobile data network. It helps the subscriber maintain contact and communicate with the data network as he or she moves throughout the network. In terms of a market potential, a PDSN can support between 12,000 and 15,000 subscribers. However, a decentralized network is also a mitigating factor (subscribers are spread over a large area, which requires more PDSNs). For a nationwide CDMA carrier such as Verizon, we believe it could eventually require as many as 2,000 PDSNs.

In the 2G world, the Interworking Function (IWF) serves as a "quasi" PDSN by offering 2G CDMA subscribers access to the data world. In some vendors' solutions, the IWF can be migrated to a 3G PDSN through a software upgrade.

In a mobile IP environment, an IP address is assigned to each mobile device. When the device moves through the network, the data side of the network needs to keep track of

FIGURE 13

CDMA Core Network Diagram

Source: DBAB

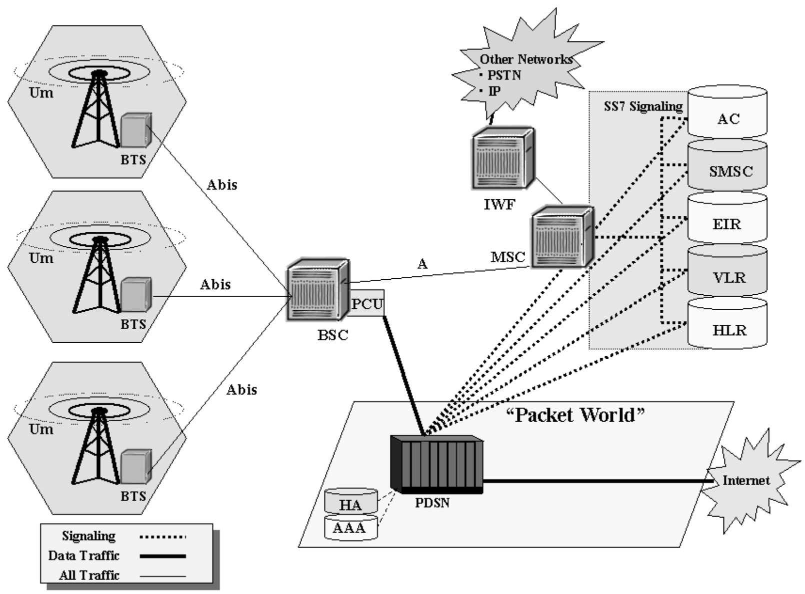

the subscriber's location in order to send data packets, just as the MSC/HLR/VLR keeps track of the subscriber for 2G voice calls. The following few paragraphs illustrate how IP mobility is maintained in a network.

When a mobile device wishes to transmit data, it sends a signal to the packet control function (PCF), which is located at the BSC. The BSC, in turn, sends this information to the PDSN, and a virtual link is established between the mobile device and the PDSN. Before the network accepts the device, the PDSN acts like a nightclub bouncer and checks out the device's credentials, or user information, with the AAA (authentication, authorization, and accounting) server. The AAA server verifies the subscriber's membership to the club, provides a mobile IP address for the session, and works with the PSDN to keep track of the tab (billing information).

As long as the device remains within the support area of the PDSN, which can be multiple BSCs, everything remains the same. However, things change if the device ventures into another PDSN's domain. If a mobile device enters an area outside of its home area (different PDSN), it still wants to continue to send and receive messages using its original IP address. In order for this to happen, the mobile device goes through the same verification process with the PDSN, in this instance called a foreign agent (FA). The FA (PDSN) passes the request to the AAA server. The AAA server passes this request to the originally assigned PDSN, or home agent (HA), to inform the HA of the devices new location. The AAA and the two PDSNs (HA and FA) then work together to set up a direct link between the two PDSNs. Once this link is set up, the HA (original PDSN) forwards data traffic to the FA (new PDSN), thus creating a true mobile IP environment.

Forget the Higher Data Rates—3G Is All About Added Voice Capacity

Despite all of the excitement over increased data rates, multimedia applications, and other feature-rich services, we believe one of the biggest drivers behind 3G is added voice capacity. Today's networks are reaching, or have already reached, saturation levels. As a result, more calls are being dropped and subscribers are increasingly having difficulty securing an open voice channel to place a mobile telephone call.[10]

Further, with worldwide market penetration at an estimated 14.7 percent by the end of 2001 and more than one-half of the world's population yet to make a phone call, we believe that lower-cost (per call) and higher-capacity mobile networks are perhaps the answer.

In order to determine how much additional capacity is added by next-generation technologies, one needs to take into consideration two facets. First, how many users can an individual RF channel can support. For example, one of the biggest drawbacks of AMPS is that it is an analog technology that does not allow more than one caller to use an RF channel at any given time. The second consideration is frequency reuse, or how often a carrier can reuse a frequency among cell sectors. One of the biggest advantages of WCDMA and CDMA is that they have a frequency reuse factor of one, meaning that the same frequency can be reused at adjacent cell sectors.

We present a chart and table (see *Figures 14* and *15*) that will probably cause more controversy than any other chart or comment we make in this report (trust us, it already has). The table illustrates how voice capacity varies across the different 1G, 2G and 3G technologies. After discussions with several OEMs and operators, we elected to present a range of values for some technologies. We took this approach because there are myriad factors that can swing the actual number in either direction.

With respect to 2G solutions such as GSM and IS–95, new voice codec algorithms can greatly improve voice capacity over initial technologies. For IS–95, this enhancement is deployed today. For GSM, the added capacity feature, called an AMR speech codec, will be available next year. Nokia also pointed out another feature, called dynamic frequency and channel assignment, which they believe can boost GSM capacity even further (by about 40 percent), although it is not being used today, despite its multi-year availability.

From our point of view, capacity claims on paper—be it CDMA, WCDMA, or GSM—do not necessarily equate to capacity improvements in networks. The reason is that operators—or for that matter, subscribers—may not be willing to accept potentially degraded voice quality. In CDMA and WCDMA networks, potential interference problems, both from in-cell callers and adjacent-cell callers, is also a factor.

For example, SK Telecom (using 1xRTT) informed us that it could support 40 voice callers per carrier, but that the operator is limiting the number to 38 callers to maintain good voice quality. At the same time, Nokia believes that 1xRTT will only result in 40 percent improvement over IS–95, which equates to 25 users per 1.25 MHz carrier. We support the higher estimate, although we recognize that other networks may not achieve the same results.

In CDMA–based networks, the best way to increase the number of simultaneous subscribers is to adjust each mobile station's transmitter power so that the signal-to-interference ratio received at the base station is at the minimal acceptable level. Otherwise, subscribers sitting on top of a base station could easily drown out far away subscribers. Thus, one of the main reasons for the increase of subscribers from cdmaOne to CDMA2000 is improved power control. Another mitigating factor is base-station density—a factor that is subject to the whims of the wireless operator.[11] In simple terms, the denser you build a CDMA network, the higher its capacity (densification).[12] This is a key advantage of this technology and is particularly true for high-density population cities such as Beijing or New Delhi.

Factors to consider when trying to interpret *Figure 14* and *15*:

- 1xEV-DO is a separate data-only channel that does not affect voice capacity.
- GPRS and EDGE traffic are data-only. Although more efficient data transfers free up channel capacity, this phenomenon is difficult to measure, and we exclude it from our calculations.
- Comparisons between CDMA 1xRTT and WCDMA are natural to make, but we emphasize that too many variables could influence actual results, thus limiting the value of a side-by-side comparison.
- Recent GSM enhancements meaningfully increase GSM capacity but have yet to be deployed.
- The transition from commercial CDMA IS–95 to CDMA 1xRTT occurs by simply upgrading an existing IS–95 network and has resulted in a capacity increase of 75 percent in a deployed network. Results for other carriers could vary and depend upon several factors, including adjacent cell interference and cell-site deployments.
- We favor the higher estimates for CDMA IS–95 and 1xRTT, although we do recognize that results in other networks could vary.

Carriers deploying WCDMA have to deploy an entirely new network to achieve the cited voice capacity. One artifact of deploying two networks (GSM, WCDMA) is that GSM capacity is additive to the WCDMA capacity, but at a high cost.

FIGURE 14

Voice Capacity per 5 MHz of Spectrum

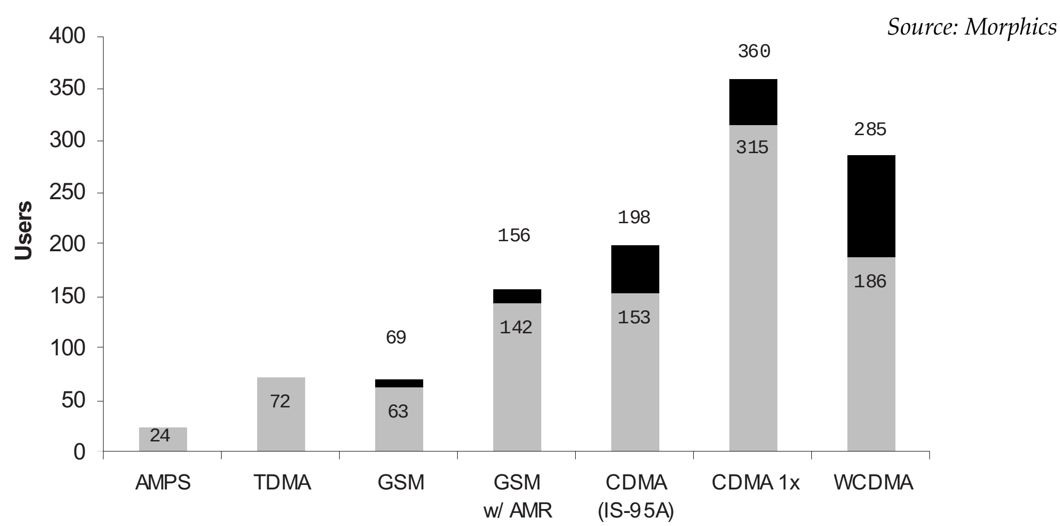

FIGURE 15

Voice-Capacity Assumptions

	AMPS	TDMA	GSM		CDMA (IS-95A)	CDMA 1xRTT	WCDMA
			BCCH[1]	Voice			
BW/Carrier KHz	30	30	200		1,250	1,250	5,000
Users/Carrier	1	3	7	8	17-22	35-40	62-95
Frequency Reuse	7/21	7/21	12	7	1	1	1
Carriers/Sector/5MHz	8	8	1	2	3	3	1
#Users/Sector/5MHz	8	24	21-52[2,3]		51-66 [4]	105-120 [4]	62-95 [5]

Source: DBAB

1) BCCH is a control channel involves using a combination of control and voice voice channels to increase network capacity. We present the assumptions for voice and control channels (e.g., user/carrier, freq reuse, carriers/sector/5MHz).
*2) Lower limit for GSM assumes 21 users per carrier. Upper limit is based upon a frequency reuse factor of 6 which gives 23 voice users versus 21 users (69/3 = 23) and AMR, which is claimed to increase capacity by 125%. (23*2.25 = 52). Available next year.*
3) DCA (not currently used, but talked about for number of years) is claimed by Nokia to increase GSM capacity by 40%. This increase is not reflected in our numbers, but could increase GSM capacity.
4) QCOM, SK Telecom and Nokia provided IS-95A and 1xRTT numbers. QCOM and SKT support the upper limit for IS-95 and 1xRTT. Nokia believes the lower estimates for IS-95A are more realistic and disputes the numbers we quote for 1xRTT. We support the upper limit range for CDMA IS-95A and 1xRTT since they are based upon actual network performance, although results could vary.
5) A more efficient voice codec could increase WCDMA capacity beyond 65 users. QCOM and NOK are not in agreement with the upper limit estimate which was made by Nokia. Given the wide assortment of variables that reflect actual results, a direct comparison with CDMA 1xRTT is difficult.

3G Handset Challenges

DSPs and Microprocessors—Better Eat Your Wheaties

Back in the good old days of first-generation technology, DSPs operated at 5 volts and had to handle only 10 million instructions per second (MIPS),[13] primarily because AMPS is an analog system. Then along came digital technology, and the MIPS requirement quickly jumped by a factor of at least 2x to 4x. As wireless systems move from second-generation voice-centric circuit-switched networks to third-generation voice-and-data, packet-switched networks, the required processing jumps dramatically by as much as 1,000x.

Next-generation phones will need to support multiple bands (e.g., 2000 MHz, 1800 MHz), multiple modes (WCDMA, GSM), multiple modulation techniques (GMSK, QPSK) and handle multiple time slots with GPRS and EDGE in order to increase the data rate. Each of these functions places an additional burden on the DSPs (yes, there is more than one), the microprocessor, plus about 1 million additional application-specific integrated circuit (ASIC) gates for WCDMA processing.[14] Motorola and Nokia both estimate that the radio channel processing requirements alone for WCDMA could reach 200 MIPS and that total MIPS requirements could reach as high as 12,000 MIPS. Although we have not seen any published numbers for MIPS, we believe similar processing power will be required for CDMA2000 (EV/DV).

The processing requirements for 3G phones are distributed across three types of processing: microprocessor (MPU), DSP, and signal-processing accelerators. The MPU supports the user interface, the communication protocol, and the applications made possible by 3G. The DSP must support increasingly complex communications requirements, as well as voice code processing. In a January 2000 IEEE article, Texas Instruments states that only 10 percent of the processing power with 3G can be accomplished on a DSP, whereas with 2G, it is 100 percent.[15] Thus, 3G phones will require additional horsepower that is nearly 10 times beyond what today's DSPs are actually capable of providing. The remaining 90 percent of physical-layer processing requires application-specific processing acceleration. *Figure 16* illustrates the distribution of 3G functionality across processing implementations.

The increased processing complexity poses a significant challenge to power consumption, time to market, and cost. OEMs and chipset vendors are striving to deliver the integration required to keep handset costs low and still maintain long battery life, but the challenges are significant. One of the biggest challenges is in the dedicated hardware accelerators. While ASICs could be used in the past, ASICs lacked flexibility, which is unacceptable in the face of shortened time to market and rapid standards evolution.

This application layer of a 3G platform is where the general purpose MPU and DSPs must be employed. Unlike voice-only phones, the 3G applications cover a range from simple messaging to real-time video. *Figure 17* from Intel provides a nice illustration of how added functionality increases the need for more MIPS.

Battery Life—Sometimes the Energizer Bunny Cannot Come to the Rescue

In the mid 1990s, Nokia was advertising its 638 phone, which provided two hours of talk time and 26 hours of standby time. An additional extended-life battery increased

FIGURE 16

Distribution of 3G MIPS

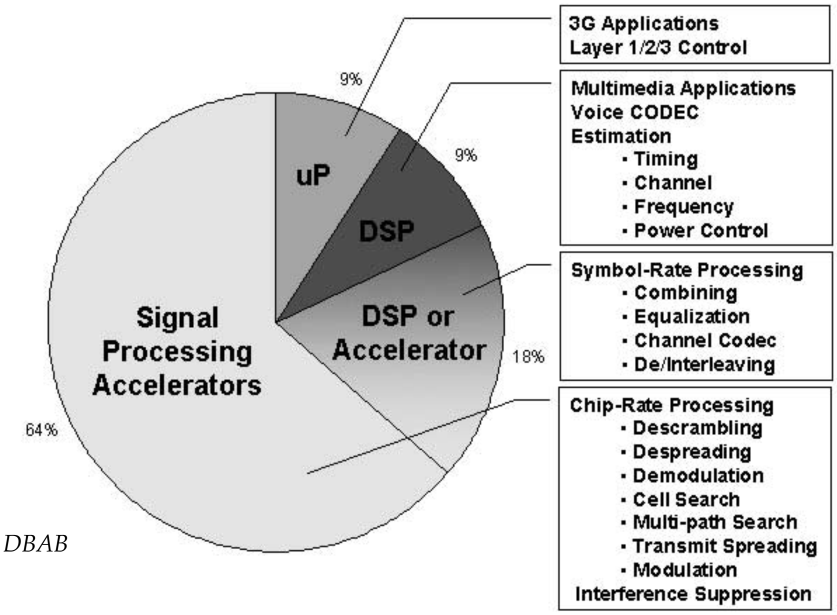

Source: DBAB

talk time to 3 hours, 20 minutes and standby time to 47 hours. Jumping ahead to more modern times, Ericsson claims its R520 phone provides 25 hours of talk time and 715 hours of standby time. To quote Bob Dylan, "the times, they are a-changing."

Battery life, until the next recharging, is determined by its energy capacity, which is expressed as a multiple of watt hours (Wh). If we assume two batteries operate at equal voltages, then the battery with the greater milliamperes per hour (mAh) rating should have a longer battery life. For example, Nokia offers an Ultralife Polymer battery that has 2,300 mAh, although most batteries have a capacity between 600 and 800 mAh. Of course, the other mitigating factor is how efficient that phone consumes power, both in standby and in talk mode. Today's premium phones consume between 2 to 4 mA in standby mode and 100 to 150 mA in talk mode.[16] We did the math, and that means a premium phone that consumes 3 mAh with a 750 mAh battery should have a battery life of 6 hours (750/125 = 6) in talk mode and 250 hours (750/125 = 6) in standby mode. Thus, the increased current required by faster DSPs and microprocessors is an important feature in addition to typical performance attributes such as processor power and memory capacity.

FIGURE 17

Application Drive MIPS

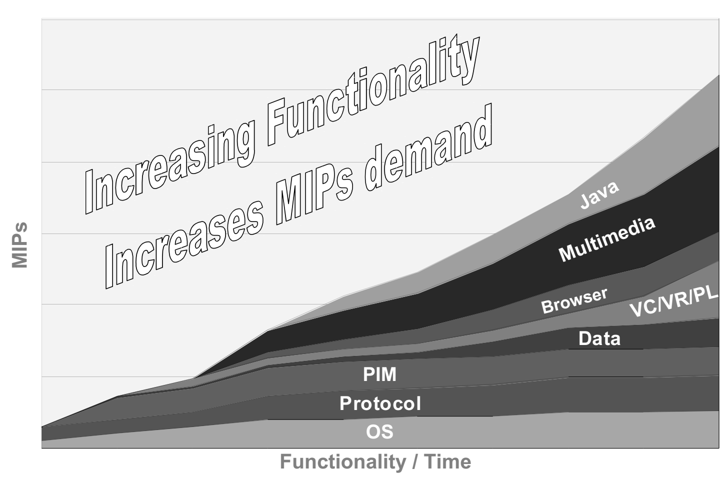

So far we have avoided using a physics equation in our discussion on 3G, but we are going to have to make an exception at this point.

Power = #Gates x Capacitance x Voltage^2 x Frequency

In order to preserve battery life, the handset needs to conserve power, sort of like Californians need to turn off the lights in unoccupied rooms in order to prevent rolling blackouts. All things being equal, a faster running processor/DSP drains the battery more quickly than its predecessor does (the frequency increases). However, the saving grace is the voltage term, which is squared in the preceding equation. If the operating voltage is reduced by a factor of 2, power consumption falls by a factor of 4 ($2^2 = 4$). Perhaps now it is apparent why companies such as Intel stress power and operating voltage requirements in addition to raw processing power.

There have been a number of battery technologies that have evolved over the years. Today, the two dominant technologies are lithium-ion (LiON) and nickel metal-hydrides (NiMh), with LiON rapidly winning the battle due to its ability to maintain its capacity after frequent recharging. It also appears that next-generation batteries will be based around an enhanced lithium technology, such as lithium-polymer, which reduces the size and improves the talk and standby times.

Don't Forget Memory Requirements

In addition to the need for extra horsepower and battery life, memory requirements with 3G also increase. It is estimated that the memory needs of a DSP read-only memory (ROM) for an EDGE handset could increase by a factor of at least five times and that DSP random access memory (RAM) will at least double. These estimates are even worse for WCDMA. The memory requirements for the microprocessor increase as well. It is estimated that Flash memory will have to increase from 16–32 Mbits to 64 Mbits and that SRAM will have to increase from around 2–4 Mbits to as high as 32 Mbits.[17]

Our Handset Market Forecast

Despite all of the challenges that exist, continued push-outs of WCDMA commercial networks and questions regarding customer usage of mobile data, we still believe 3G handsets will eventually come to market, if nothing more than as a means of offloading voice traffic from today's capacity-constrained networks.

Our forecast divides the world into five regions: North America, Europe, Latin America, Asia-Pacific, and Africa/Middle East. For each region, we project on an annual basis the total population and installed mobile subscriber base for each region, which we believe will increase each year as mobile phones permeate our society. We currently forecast a continuous annual growth rate of –9.0% (2001 to 2006), suggesting that growth in the industry will have to come from replacement sales.

Luckily for the industry, handsets are not a durable good and are replaced on a periodic basis. Since the replacement rate varies by technology, region and year, it would be too cumbersome to present a six-year forecast. However, we can summarize our assumptions by stating that we believe replacement handsets will represent 31 percent of the handset market in 2002 and reach 86 percent of the handset market in 2006.

As *Figure 18* illustrates, we currently estimate 389 million unit sales in 2001, 440 million unit sales in 2002, and a CAGR (2001–2006) of 14.0 percent. We also currently believe that WCDMA unit sales will grow from a 0 percent market-share to represent 24.2 percent of the handset market in 2006. We also believe that CDMA unit sales will account for 24.5 percent of the market in 2006, with a majority of these handsets being next-generation handsets.

Our Mobile Infrastructure Market Forecast

We estimate future mobile network capital expenditures (CAPEX), based primarily on new subscriber growth and increased minutes of usage (MOUs), from both voice and data. The increased CAPEX per each additional subscriber varies by region and by technology and should decrease over time. For example, we estimate that a GSM carrier will spend $137 for each new subscriber that it adds to its network in 2001 and that a TDMA carrier will spend $98. We estimate this amount will fall to $120 and $86 in 2002, respectively. The results of our 50,000+ cell Microsoft Excel model are presented in *Figure 19*.

FIGURE 18

Our Handset Industry Forecast by Technology

	2000	2001E	2002E	2003E	2004E	2005E	2006E	CAGR 2001-2006
Handset Unit Shipments (000's)								
Analog	14,046	1,843	306	-	-	-	-	-100.0%
GSM	266,860	238,664	272,075	295,231	297,948	311,131	313,712	5.6%
WCDMA	-	0	2,317	14,850	56,572	110,879	181,272	1260.8%
CDMA	51,698	69,988	86,486	112,479	134,913	158,784	183,410	21.2%
TDMA	42,369	43,719	46,583	49,674	51,348	52,125	52,292	3.6%
iDEN	4,005	4,394	4,690	5,019	4,548	4,851	4,660	1.2%
PDC	27,925	30,390	27,438	26,236	24,022	19,959	13,135	-15.4%
Total Unit Shipments	**406,902**	**388,999**	**439,894**	**503,490**	**569,351**	**657,729**	**748,482**	**14.0%**

FIGURE 19

Our Mobile Infrastructure Forecast by Technology

By Technology (mils)	2000	2001E	2002E	2003E	2004E	2005E	2006E	CAGR 2001-2006
Analog	578	223	96	-	-	-	-	-100.0%
GSM	28,323	23,641	21,387	16,624	12,893	10,888	8,995	-17.6%
WCDMA	-	695	4,084	13,614	24,017	33,059	39,257	124.1%
CDMA	10,757	11,091	11,938	10,840	10,227	9,433	8,841	-4.4%
TDMA/UWC	3,964	3,853	3,343	2,710	2,189	1,780	1,392	-18.4%
iDEN	470	433	362	322	210	179	98	-25.7%
PDC	2,170	1,752	1,200	867	658	422	171	-37.2%
Totals	**46,262**	**41,687**	**42,410**	**44,977**	**50,195**	**55,761**	**58,753**	**7.1%**
Cell site installs	140,064	140,776	159,061	189,723	243,859	310,358	373,570	21.6%
Total # of cell sites	475,095	615,870	774,931	964,654	1,208,514	1,518,871	1,892,441	25.2%
Switch installs	3,602	3,607	4,008	4,557	5,605	6,940	8,191	17.8%
Total # of switches	11,948	15,556	19,564	24,122	29,727	36,667	44,858	23.6%

At first it may appear a bit dismal that the mobile infrastructure spending CAGR (2001–2006) is only 7.1 percent. However, our forecast includes decreased CAPEX for analog (goes to $0 in 2003) and 2G technologies (our CDMA forecast combines 2G and 3G, since it is difficult to separate future CDMA CAPEX between 2G and 3G installations). Since most investors focus on 3G (WCDMA) infrastructure spending, we point out that the CAGR for WCDMA is 124 pecent, growing from only $695 thousand in 2001 to $39.2 billion (66 percent of total annual CAPEX) in 2006.

Conclusions

- We currently estimate handset sales of 389 million in 2001 and 440 million in 2002, with WCDMA handset sales not taking off until late 2002 at the earliest and more likely 2003.
- We currently estimate mobile infrastructure CAPEX spending of $41.7 billion in 2001 and $42.4 billion in 2002. While we forecast a mobile infrastructure CAPEX CAGR of only 7.1 percent (2001–2006), we also forecast WCDMA infrastructure CAPEX growth really taking off in 2003 with a CAGR (2001–2006) of 124.1 percent.
- Based upon discussions with industry experts, we believe that WCDMA will eventually be successful. Currently, the standard is relatively immature and the technology is going through the same growing pains that the CDMA IS-–95 standard went through when it was first deployed.
- We believe 3G infrastructure spending, at least for the near-term, will be driven by the need for additional voice capacity. CDMA2000 and WCDMA technologies both provide the necessary voice capacity that operators so desperately need. The question becomes, when will carriers "cave in" and begin spending CAPEX dollars.

Note: This paper was condensed from a 200-page industry report that was published in September 2001.

Notes

1. In 3G WCDMA networks, the BSC is called a radio network controller (RNC).
2. A T1 line supports 1.544 Mbps, while a T3 line handles 44.7 Mbps. Their European counterparts are the E1 and E3 lines, respectively.
3. As we discuss later in this report, nearly all 3G networks will be based upon a form of CDMA, CDMA2000, or WCDMA.
4. GSM originally stood for Groupe Special Mobile from approximately 1984 to 1992.
5. Verizon and Cingular currently offer similar plans, starting at $2.99 per month for up to 100 messages. Additional messages cost $0.10 each (Verizon only charges $0.02 to receive a message).
6. 3G (UMTS) is comprised of a series of releases, beginning with Release '99.
7. Multiple received signals generally result in poorer reception due to an effect call multi-path. However, a CDMA phone uses a RAKE receiver that can sort through and combine multiple signals, just as a yard rake can group together individual leaves into a larger pile.
8. Comments in this section are based upon discussions with Ericsson and the book *WCDMA for UMTS*, edited by Harri Holma and Antti Toskala of Nokia.
9. In our GSM and WCDMA sections, we introduced a number of interfaces (A, Abis, etc). The ANSI–41 protocols have essentially the same naming scheme and the same non-technical description, so we elected to exclude them from our CDMA (ANSI–41) section.
10. Ironically enough, at the recent Cellular Telecommunications & Internet Association (CTIA) conference in Las Vegas, it was nearly impossible to place a call through our mobile carrier's network.
11. Our IS–95 and 1xRTT estimates are based upon information obtained from vendors and carriers deploying commercial CDMA IS–95 and 1xRTT networks. These estimates could vary in other carrier's networks. WCDMA estimates were taken from the book *WCDMA for UMTS* by Henri Holma and Antti Toskala of Nokia.
12. Densification is a term we invented, but we hope it gets our point across.
13. MIPS is sometimes referred to as Meaningless Indicator of Performance, since it only provides a rough estimate for measuring performance.
14. The primary difference between a DSP and a microprocessor is that a DSP is designed for tasks that require repetition and are numerically intensive (e.g., voice encoding). A microprocessor is more robust and flexible and is designed for control-oriented applications.
15. Texas Instruments, "DSP–Based Architectures for Mobile Communications: Past, Present and Future," *IEEE Communications Magazine*, January 2000.
16. Paul OuYang, *Communication Systems Design*, March 18, 2001.
17. Paul OuYang, *Communication Systems Design Magazine*.

Unlock the Value of Convergent Wireless Services: Pricing Models to Achieve Optimal Profitability

James Morehead
Director, Wireless Market Development
Portal Software, Inc.

Introduction: The Wireless Industry in 2001

"Value pricing appears in other contexts. People pay more for orchestra seats than for balcony seats; for Saturday night performances than weekday performances; for the services of more skillful doctors or consultants. … Smart marketers will bundle the product with additional benefits and price the total offering."

Philip Kotler, Kellogg Graduate School of Management, Northwestern University

"The business model based on paid content is simple and transparent—content providers know what they want and what revenues they will get. Users will pay for content that has a high perceived value, that is personalized and delivered in a format that suits their device of choice."

Eden Zoller, Ovum Senior Analyst, August 2001

"Most carriers' billing systems are not advanced enough to split revenue based on customer minutes, and then send revenue back to content providers."

Unstrung, August 2001

"'Content will soon be developed and provided on a usage basis to different market sectors,' predicts Greg Howard, principal analyst and founder of High Tech Resource Consulting Group in San Andreas, California. 'However, a significant hurdle in developing these new business models has been the lack of content usage-based billing solutions.'"

Global Telephony, June 2001

"'We see that there will be traction for content-based billing in two or three years; once … the bandwidth is there with DSL and GPRS,' says Jason Briggs, senior analyst in the Yankee Group's Billing and Payment Applications Strategies division. 'But this is the time where the winners will get the mind share.'"

Billing World, April 2001

"'In most cases operators are still assessing market requirements through service trials, so strategies are still being redefined, but the favored model at present is to charge by the amount of data consumed,' said Yunus. 'Current billing platforms will be inadequate both in the long term and in stimulating demand right now, given the diversity and range of future data applications,' he added."

"Mobile Data Pricing – Will European Consumers Pay a Packet?" Yankee Group, August 2001

"China is also experiencing significant growth in SMS traffic, according to ING Barings, after usage was initially hindered by flat-rate pricing models offered by China Mobile. 'The company later changed the charging method to a per-message basis which started stimulating SMS usage,' says Barings analyst Craig Racine."

Global Mobile, August 2001

Value-Based Pricing Overview

In the mobile industry, value-based pricing (also referred to as content-based pricing or event-based billing) is being treated as a new way of charging for services. While it is true that most mobile services to date have been charged on a per-minute or subscription basis (voice-mail, calling features), packaged products have long been priced based on value.

Before applying value-based pricing concepts to mobile services, it is useful to consider value-based pricing successes in other markets. The first of several case studies used in this paper is bottled water (refer to *Figure 1*). The content delivered in bottled water is arguably "free" (while in some cases the source it a distant mountain spring, in many cases the content is filtered tap water). There is added value; however, the water is delivered cold and in a portable package—and people are prepared to pay extra for this value—a "convenience premium." There is also perceived value—savvy marketing positions bottled water as "safer" than tap water—and reinforces that water is "healthy" and that "water is life".

FIGURE 1

The Power of Perceived Value

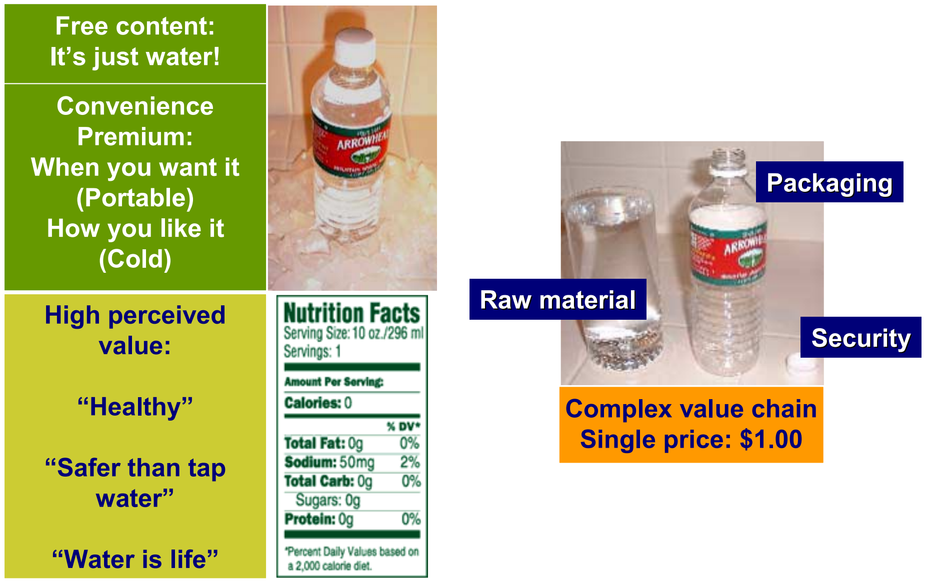

This simple example parallels mobile services—the core value of a mobile phone is its convenience (the "freedom to socialize" with friends and family from any location; the ability to conduct business while on the move). In most countries, there is still a convenience premium paid for mobile phone service over fixed lined service (although this is shrinking over time).

Even more striking is the complex value chain that is required to deliver an ice-cold bottle of water into the hands of a thirsty consumer—elements of the value chain include raw materials, packaging, physical distribution, and even security mechanisms (the tamper-proof cap). Despite the complexity of the value chain, bottled water is sold for a single, simple price—and the price is clearly labeled.

A second example is outlined in *Figure 2*. Here, consumer products are used to compare value-based pricing with "weight-based" pricing. Most consumer products are value-priced—the "weight" of the product is not related to the price, and "raw material" costs are hidden from the purchaser. Even "commodity" products can be branded and their value increased (the bottled water example).

Let's consider for a moment a world where all products are sold based on their weight. In this world, a single price per pound is set for all products (in the example in *Figure 2*, $4.00 per pound). Products such as baby formula with high perceived value are under priced. "Heavy" products such as baking mix benefit. And the cotton ball industry wouldn't have a business case (or, the cotton ball industry would focus on genetically engineering "heavy cotton"). The value chain is encouraged to create "heavy" products and discouraged from creating "light" products.

In summary, people (consumers and business people) accept per-product/per-service pricing and "points of value" as long as the pricing is clear and transparent. Value-based pricing models rewards innovation rather than just "heavy" products.

Fundamentally, why are mobile services any different? This paper argues that mobile services are ultimately no different than the simple examples shown and submits that charging for the value of content and services, not the bytes (raw materials) used to create services, will enable the mobile industry value chain to maximize the profitability of convergent wireless services.

Value-Based Pricing Applied to Convergent Wireless Services

For illustrative purposes, a wireless digital picture sharing service is referred to throughout this paper (such a service is available through both J-Phone and DoCoMo in Japan with multimedia handsets). With this service, users are able to take a picture with a digital camera (integrated into the handset or connected via a cable or Bluetooth), and then send the picture wirelessly to another handset.

The service has value—it is more convenient to send a digital picture directly from a portable device than to synchronize the device to a computer for later transmission. The service is also immediate—for consumers, this makes the service "fun," and, in the case of business applications, provides a competitive edge (e.g., a photo journalist trying to secure a "scoop" at a crime scene).

The service described has substitutes—the end user (for example, a tourist) could take a picture with a digital camera and later upload the pictures to a computer for sending. Or that end user could buy a postcard and stamp (perhaps $0.60 combined). In both cases, the substitute is less convenient and immediate (and in this example, less personal).

FIGURE 2

Decoupling Product Cost from Product Value

Having concluded that the service is of value, how should the service be priced? Three key options are subsequently outlined.

Transport Model (Transport Charge Only)

The simplest pricing model for mobile service providers is just to count the number of bytes transmitted and apply a price per byte. This model scales along with traffic but does not provide any differentiation between the types of services delivered. Low byte volume/high value content generates less revenue than high byte volume/low value content. Just as was the case in the introductory examples (with consumer goods sold by weight), the transport model rewards applications that consume more network capacity and undervalues innovative services that minimize network usage.

Further, the transport model makes "sender pays" pricing models difficult to implement because the type of service is more difficult to determine from a network-level byte stream (it is difficult to "sniff" out services when there are potentially hundreds of content providers). Services such as the digital picture-sharing example will likely be more widely adopted in a "sender pays" model (the wide adoption of core voice service in Europe is partially attributed to "calling party pays" pricing models).

The transport model also makes it difficult to decouple the raw material cost (bytes) from the service value. For example, as compression technology improves, the transport model would automatically pass savings along to the end user, since it is network level activity, not service-/content-level activity that is monitored. The price drop may not be necessary if the compression technology maintains the service quality. With the transport model, it is also difficult to implement content-specific loyalty programs (such as send five digital pictures, get one free).

The transport model is the least intuitive and requires the end user to understand how the service "works" in order to predict the price (e.g., performing a mental calculation: the number of bytes required for the service multiplied by the price per byte). And when the invoice arrives (for a postpaid subscriber), the charges will not be itemized by service but simply shown as a block of bytes. Prepaid users may be afraid to use services under a transport model because they won't be sure how their balance will be impacted.

It is important to note that there are cases where the transport model is appropriate—corporate intranet access, for example, where an enterprise has purchased access in bulk. In general, however, the transport model makes it difficult to align pricing models with service value and to built service-specific remittance agreements with content partners.

Hybrid Model (Content Charge + Transport Charge)

In the hybrid model, there is a charge for both the picture sent and the transport. This requires the mobile service provider to track content events and transport events but not correlate the two. In this model, there is the ability to charge for specific points of value. Because the content charge component is not affected by compression technology improvements, the impact of compression technology on revenues is less severe. The content component enables a partial "sender pays" model so that the receiver is only charged for the transport.

The pricing model, however, is still complex and requires the end user to make two decisions: "Is the content charge good value? What is the additional transport charge?" As with the transport model, the end user will likely not know how to calculate the transport charge, which may delay service adoption. The hybrid model is not favorable to prepaid because of the split charges for a single service.

Content Model (Content Charge Only)

The most intuitive model is charging solely for the content—in this case, the digital picture (e.g., a charge per picture sent). Charging solely for the content simplifies "sender pays" pricing models (since user- and service-specific event information is received from the content server rather than the network). The content model also makes loyalty programs easier to implement (send five digital pictures, get one free).

To simplify the pricing model for the end user, the value chain must absorb complexity. In this example, "absorbing complexity" means not charging for bytes that are correlated to services charged on a content basis, and processing event information received from a third-party content provider.

Table 1 provides a summary of the aforementioned models.

As summarized graphically in *Figure 3*, the impact of compression technology on service revenues and margins is dramatic in the transport model and hybrid model. Only the content model is unaffected by the introduction of compression.

The Advantage of Convergent Services

The digital picture example can be further enhanced when converged with voice services. For example, consider a "Friends and Family Plus" plan where digital pictures sent between friends and family members are coupled with free voice minutes (e.g., two voice minutes per digital picture sent). Voice minutes are used to encourage the use of a higher margin service—sending digital pictures. The example shown in *Figure 4* demonstrates two subscribers in a friends and family group sharing digital pictures, and then making free phone calls to talk about the pictures.

To connect free minutes to services, the hybrid or content model is required. Assuming the content model ($0.75 per digital picture) and two free minutes per picture sent, the resulting revenue is $1.50. The transport cost is four minutes

TABLE 1

Pricing Model	Revenue Generated[1]	Mobile Service Provider View	End-User View
Transport Model ($5.00 per MB)	$0.15 sender – transport charge $0.15 receiver – transport charge	• Revenue decreases with the introduction of compression technology • Loyalty programs difficult to implement • Difficult to differentiate • Difficult to implement "sender pays" models	• Non-intuitive; service may be perceived as more expensive than it really is • Forcing receiver to pay could hamper service adoption • Invoice does not show content-usage detail; repudiation risk (sender and receiver)
Hybrid Model ($0.30 per Digital Picture and $5.00 per MB)	$0.30 sender – content charge $0.15 sender – transport charge $0.15 receiver – transport charge	• Revenue decreases with the introduction of compression technology • Difficult to implement "sender pays" models • Enables loyalty programs and content-specific bundles	• Non-intuitive; service value decision is based on two prices • Forcing receiver to pay could hamper service adoption • Invoice does not show content-usage detail; repudiation risk (receiver) • Separate charges complicates prepaid models
Content Model ($0.75 per Digital Picture)	$0.75 sender – transport charge No charge to the receiver	• Enables loyalty programs and content-specific bundles • Compression does not impact revenues; rather, compression increases profits by reducing transport costs • Well suited to prepaid subscribers	• Intuitive; service and pricing model are matched; pricing is transparent • "Sender pays" model removes barrier for receiving pictures • Invoice is content-specific and reduces repudiation risk • Prepaid and postpaid pricing models can be the same

FIGURE 3

Impact of Compression for Transport, Hybrid, and Content Models

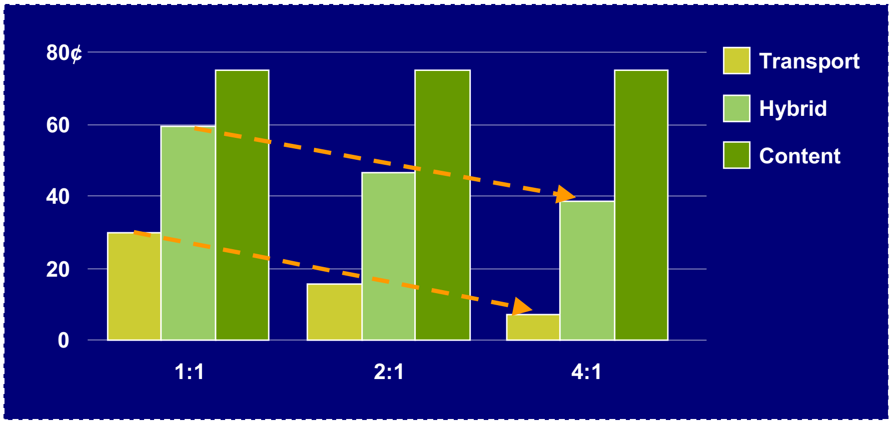

of voice (the free minutes) and a voice minute equivalent of less than two minutes for the data transferred (approximately 120KB)[2]. The convergent offer not only encourages use of a higher-margin service, but also builds loyalty within a friends and family group.

If the value chain cannot capture the full value of content and services, the revenue stream will be insufficient and the value chain collapses. Without an incentive to create interesting content and services, the value chain cannot exist.

The convergent example discussed here is similar to the strategy used at supermarkets—the milk (a high-volume, commodity-priced or "everyday low price" product) is placed at the back of the store, whereas high-margin products are place along aisles leading to the milk (the "margin gauntlet"). As mobile service providers expand their service and content offerings, there will be increasing opportunities to cross promote, cross discount, and bundle services. Charging for content rather than charging for bytes transmitted greatly increases the options for marketers.

A single example has been examined in detail to draw out the advantages of value-based pricing and convergent services. There will, of course, be a variety of pricing models employed—flat-rate services, usage-based services (metered on content and transport volume), and subscription services. The most successful mobile service providers will build pricing models to match market-driven requirements, not billing system limitations.

The evolution of pricing models in the mobile industry as voice, data, and content converge and competition increases is inevitable—one need only look at other indus-

FIGURE 4

Converging Voice and Data Services

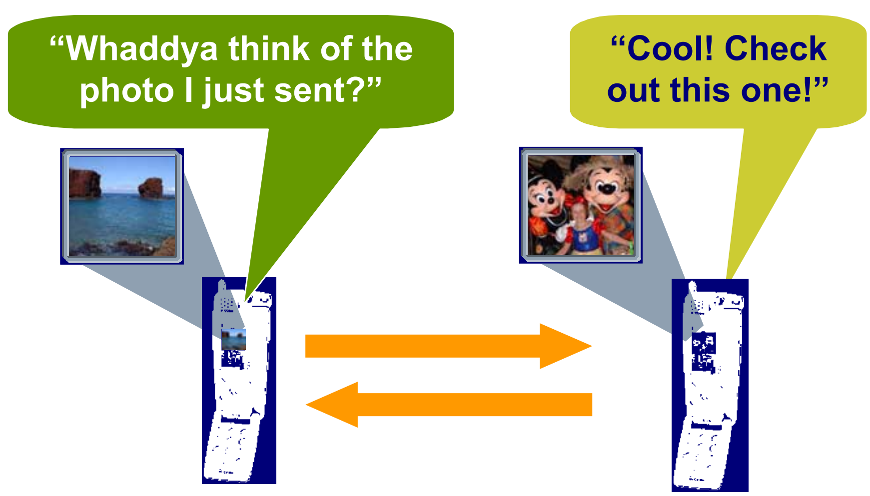

tries to see how pricing evolves in the face of competition (see *Figure 5*).

The key question then is how to implement value-based pricing models effectively. This is addressed in the next section.

Building a Content-Enabling Environment

To enable the digital picture-sharing service described, a new relationship between the mobile service provider and the content provider is required. The content provider will usually be in a better position to specify the detailed nature of a transaction and will hold the copyrights for branded content. In addition to providing the transport, the mobile service provider will usually be in a better position to obtain payment from the end user (the *only* way in the case of a prepaid subscriber that does not want multiple balances with multiple parties). The mobile service provider is also an excellent source of location information to enhance content and services.

Consider the digital picture example. The mobile service provider is able to detect the bytes going across its network, but unless it has chosen to host the service, the mobile service provider may not be able to track the number of digital pictures sent and to whom they are sent (and if the digital pictures are encrypted for privacy, tracking transactions at the network level will be virtually impossible). The content provider is able to track the number of digital pictures sent—even if encrypted—and provide detailed information regarding both the sender and receiver but may not be able to entice the end user to set up a separate account for billing (or a separate prepaid balance).

Further, if the content provider chooses to bypass the mobile service provider, the end user will pay the transport directly, resulting in the hybrid model. As discussed earlier, this is a less intuitive pricing model that may impact the rate of service adoption. When the mobile service provider is bypassed, the benefits of convergence are lost (cross-product discounting, bundles, and convergent loyalty programs), which hurts both parties.

In the digital picture example, if the mobile service provider and content provider work together, the full value of the service is "unlocked." The content provider provides usage details about the digital picture service, and the mobile service provider bills the end user (or impacts the prepaid balance). The mobile service provider and content provider then settle the revenue received. Just like buying bottled water, the end user is shielded from the complexity of the value chain. An interconnected value chain maximizes value to the end user—and as a result maximizes revenue to the value chain. There are strong motivators for the mobile service provider and content provider to work together and form a profit maximizing value chain—it is the natural evolution of the mobile industry.

The mobile service provides a "content-enabling environment" and extends its value well beyond a simple "pipe" for bytes. The content provider, by working with the mobile service provider, ensures full value for its services and enjoys the economies of scale of the mobile service provider's convergent billing environment.

This relationship is shown in *Figure 6*.

Creating New Value Chains

Next-generation billing platforms enable links between voice and data network technologies, content partners, and payment methods. Innovation will occur in the network as

FIGURE 5

The Evolution of Pricing in Competitive Markets

	Initial Pricing Options	Pricing Options in the Face of Increased Competition	

Automotive	• One price, one model, one color (Model T)	• By model
		• By bundled option package
		  • Sports option, luxury package
		• Targeted at market segments
			• Dealer incentives
			• Value-added features
			  • Tires
			  • Stereo
			  • Paint
			  • Warranty
			  • Other

Airline	• By mileage	• By class • By time of purchase • Corporate discounts • Promotional fares	• By method of purchase (Internet or CSR) • Refundable or not • By mileage points • By routing
Consumer Products	• By product	• By product • By brand • Targeted at market segments	• By location • Real-time coupons based on purchasing behavior • Promotions

FIGURE 6

A New Relationship between Mobile Service Providers and Content Providers

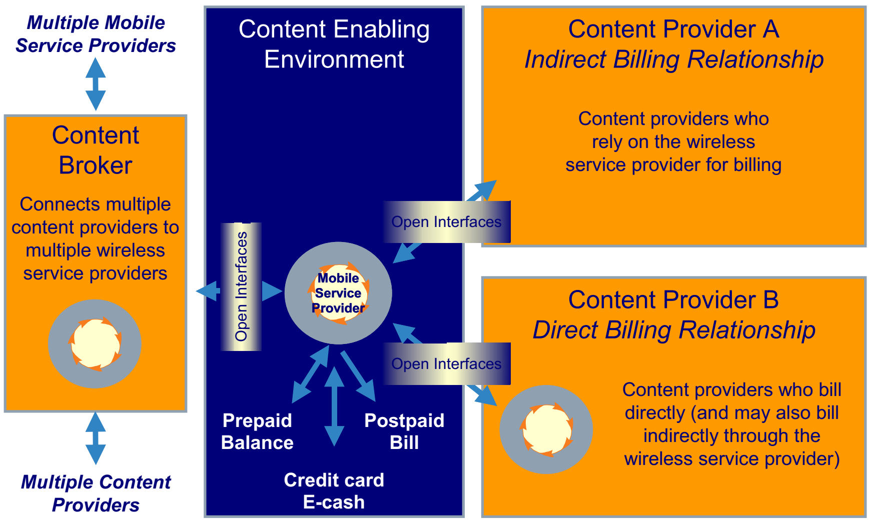

new technology enablers are deployed and outside the network as new services are created. As highlighted in *Figure 7*, traditional billing systems were custom built to interface with network and mediation elements for batch processing of voice call detail records (CDRs). Next-generation billing platforms support both batch interfaces for high-performance processing of voice CDRs and transactional real-time interfaces for prepaid services, advice-of-charge, and other applications requiring real-time processing.

The explosion of content providers outside the mobile service provider's network requires billing platforms to support open interfaces. Going forward, it will not be practical for a custom integration to be built between each of hundreds of content providers and the mobile service provider.

Customized closed systems were not designed or intended to rapidly adapt to such a constantly changing environment. The total cost of ownership (TCO) for custom-built systems increases rapidly as the number of content providers increases. In a highly competitive marketplace, a delay of even one month can have a significant impact on the profitability of a new service.

FIGURE 7

Multiple Sources of Information in a Mobile Network

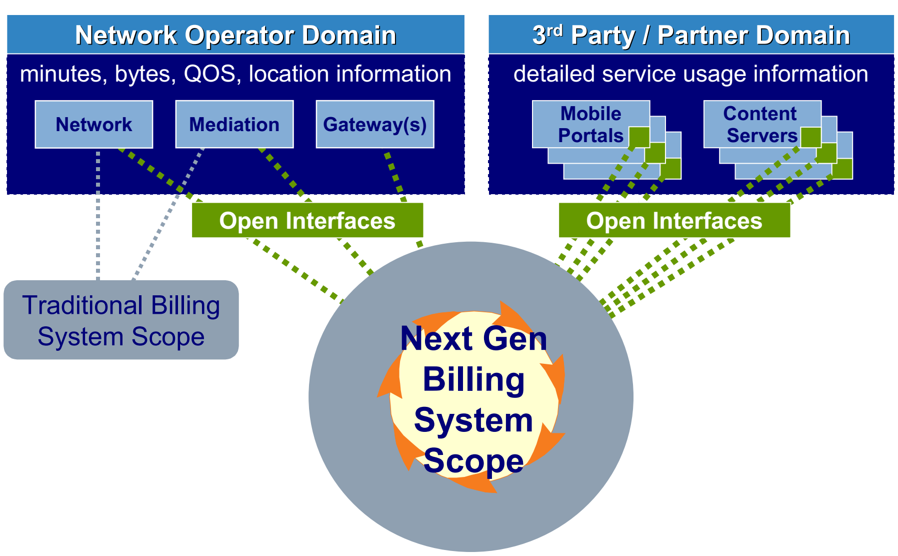

Figure 8 demonstrates the flexibility of a next-generation billing platform using the digital picture example. While the network and mediation layers provide information on the volume of bytes consumed using the mapping service, and possibly the universal resource locator (URL) of the service, the mobile portal and/or content service have very rich detail on the service consumed. This richer detail enables innovative pricing models, sponsorship, loyalty programs, and remittances, *but only if the information can be effectively delivered to the rating and billing platform.*

Conclusions

There is growing consensus in the mobile industry that value-based pricing models are required to maximize the profitability of enormous investments in 2G, 2.5G, and 3G mobile networks. Next-generation billing platforms will enable mobile service providers to match service innovation with pricing innovation. Events for rating and billing are no longer limited to network-centric sources. Rather, a next-generation billing platform enables a customer-centric approach—using sources of information that are closer to the service and the customer—enabling pricing models that are intuitive, competitive, interactive, service-specific, and, most importantly, profit-maximizing.

Notes

1. The revenue generated is based on a 30KB transfer (including overhead). In the transport model, the price is based on 30KB transmitted at $5.00 per MB. In the hybrid model, the sender pays for content and transport ($0.30 for the digital picture and $0.15 for the transport). In the content model, only the sender pays, since there is no transport charge (network activity to the content provider is "non-chargeable" but may be tracked for service profitability analysis).
2. This assumes 14.4 kbps packet data, which translates into 90KB/minute in circuit-switched terms, or less than two minutes of "voice equivalent" network capacity. 120KB represents two pictures sent, 60KB for each (30KB to send and 30KB to receive).

FIGURE 8

Multiple Sources of Information for a Digital Picture-Sharing Service

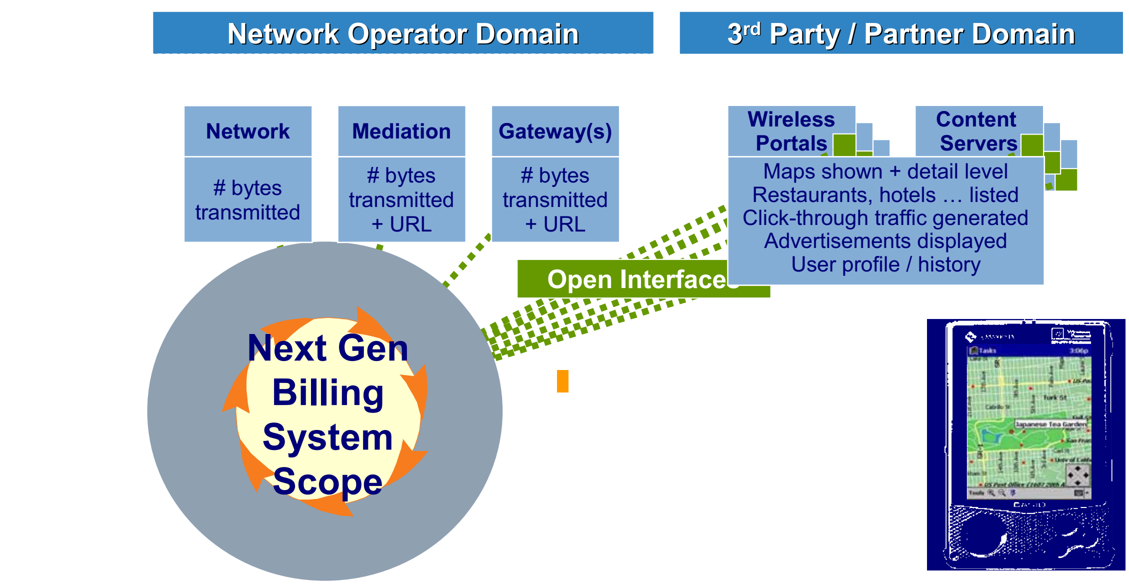

Advanced Hyper Satellite Designs for the Next Decade

Joseph N. Pelton
Professor, Institute for Applied Space Research
George Washington University

Ivan Bekey
President
Bekey Design Concepts, Inc.

Abstract

For communications satellite systems to remain relevant a decade hence, new technology and new design architectures may well be necessary. In this article, a brief description is given to four "dramatically different" designs of future "hyper satellites" that could be deployed in the 2010 to 2015 time period. These totally new types of satellites architectures, once deployed, could delivery services at extremely high rates—i.e., in the terabits/second range and at transmission costs that might be as low as $0.0l per 2 gigabits of information transmitted.[1] If the designers of satellite systems do not seriously consider breakthrough designs with new types of systems architectures within the next few years, then satellite networks could lose rather than capture new market-share in future years. This would be unfortunate, because broadband Internet, interactive entertainment, mobile Internet protocol (IP) services, and other opportunities are now open to the satellite industry.

Hyper Satellite Design Concepts

For more than 35 years, the dominant trend in satellite communications has been to develop higher-capacity and higher-throughput satellites by designing larger power systems and higher-gain antennas (and thus larger-aperture antennas). This trend has also produced a corollary assumption. This is the dominant systems paradigm where moving to larger-aperture and higher-gain satellite antennas (accompanied by higher-wattage power systems) has also meant that we had to move to much higher-mass satellites that would be more expensive to launch and also bring other problems related to reliability, etc. The antennas that are planned for the next Telesat system, as well as those designed and launched on the Asian Cellular Satellite System (ACeS) and Thuraya represent this trend toward bigger, more massive, and certainly more expensive satellites (i.e., both to build and launch).

Figure 1 shows this type of extrapolation toward larger and massive satellites with higher throughput capabilities. This same graphic also includes a very important new idea. This idea is that if there were better designs with better and lower materials then in a decade or so more productive and cost-efficient satellite systems could be developed.

There are, in fact, a large number of potential new architectures that could allow the design of new satellite systems that could be at least an order of magnitude more cost-effective than "conventional satellite communications" technology, as represented by an Intelsat 9, an ACeS or a DirecTV 4S satellite. One of the more promising approaches developed by researchers in the field is the design of satellite communications antenna systems that use remarkable lightweight materials as well as advanced and highly efficient power systems so that the throughput/kilogram is 10 times less that today's conventional designs. Also, the power levels derived from the very highly focused satellite beams would provide flux densities sufficient to support tiny micro-terminals (including the possibility of wearable mobile antenna systems).

In studies currently being conducted by George Washington University's Institute for Applied Space Research (IASR), there are three different types of advanced designs under consideration. These designs (that might be called "hyper satellite designs") all achieve very high capacities (or digital throughput) while minimizing mass and systems costs and also increasing usable lifetimes. These are briefly described as follows:

- *Hyper Satellite Option One:* Implement a constellation with a number of very lightweight antennas placed far apart in space and fed by a central unit. The location of the antennas in this constellation would be maintained principally by passive means, most likely achieved by a rotating tethered system.
- *Hyper Satellite Option Two:* Implement one or, at the most, a few very large filled aperture antennas using advanced phased array antenna technologies and little or no structure, such that one to two orders of magnitude of weight and cost reduction are obtained compared to conventional antennas. One of the principal keys to this design is that the feeds to the antennas would be phased array elements or multiple illuminators that would create many narrow beams.
- *Hyper Satellite "Swarm" Option Three:* This design would implement an extremely large sparse array antenna with a large swarm of tiny elements acting in concert to form a coherent phased array without any structure. The dimensions of the antenna would be principally maintained by passive (i.e., gravitational) means.

FIGURE 1

Past Trends in the Size and Mass of Satellites and a View of the Potential of a New Type of Satellite Architecture

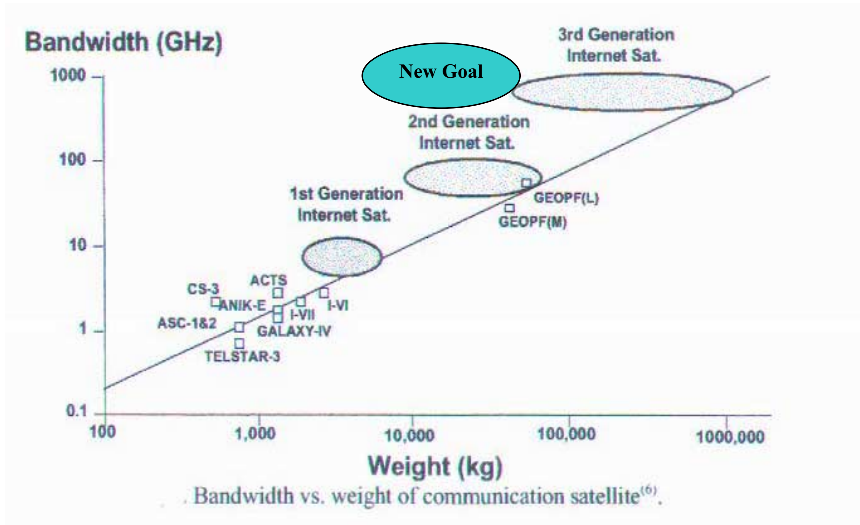

Such designs are currently being evaluated against such factors as cost, reliability, technical feasibility, complexity of deployment and operation, need for new material development and new research and development (R&D), space environment/space debris concerns, human health and safety issues, and legal or regulatory constraints.

All of these designs would also include new highly efficient "rainbow" solar cells with a very high level of "junctions" so that a much higher percentage of the power contained in infrared light and ultraviolet radiant energy would be captured and thus convert more than 50 percent of the sun's energy. Such totally new designs are not likely to be developed prior to 2010 and would not become operational until 2015. Unless key development work is carried out now, however, these highly efficient systems may not be accomplished even by those dates. Important technologies needed to support such advanced satellite platforms include those listed in *Figure 2*.

FIGURE 2

KEY TECHNOLOGIES FOR GEO PLATFORM SYSTEMS

- Advanced phased array antenna concepts
- Polyimide materials (less that one kg/square meters, including electronic components)
- Advanced fullerene-based (e.g., "Bucky ball") and other advanced tether structures for antenna deployment
  Tether structural design for forming passive constellations
- Adaptive electron-beam shaped piezoelectric membranes
- Liquid metal solid-state field effect electric propulsion units
- Broadband time-delay microcircuits that are capable of modulation
- Large phase shift capability microcircuits that can be modulated
- 10–12 bit digital samplers at gigahertz frequencies
- Integrated self-contained silicon-based picosatellites that can be deployed in swarms
- Highly efficient photovoltaic thin films (e.g., rainbow solar cells)
- Piezoelectric bimorph membranes resistant to the space environment
- High-temperature superconducting long wires
- Precision optical sensing of surface figures
- Precision metrology for position control in formation flying
- Integrated multifunction components

To illustrate what these concepts might look like in actual operational satellite systems, the following graphics are provided as a preliminary visualization of these types of advanced communications satellite systems. Concepts 1, 2, and 3 as presented in *Figures 3*, *4*, and 5 would represent increasingly more demanding new technology.

Beyond these three new rather radical "hyper satellite" design concepts, there is yet another way to devise telecommunications concepts for the future that are more than an order of magnitude more cost-effective than conventional satellite technology. This is merger optical laser technology that is being developed for intersatellite links with so-call stratospheric platforms or high-altitude platform systems (HAPSs).

High-Altitude Platform Systems Including Hybrid Satellite and Stratospheric Platform Systems

Thus, the other "new idea" is to have space communications service to HAPSs via optical links and then using radio frequency (RF) service (probably in the Ka-bands or above) to

FIGURE 3

Hyper Satellite Concept 1

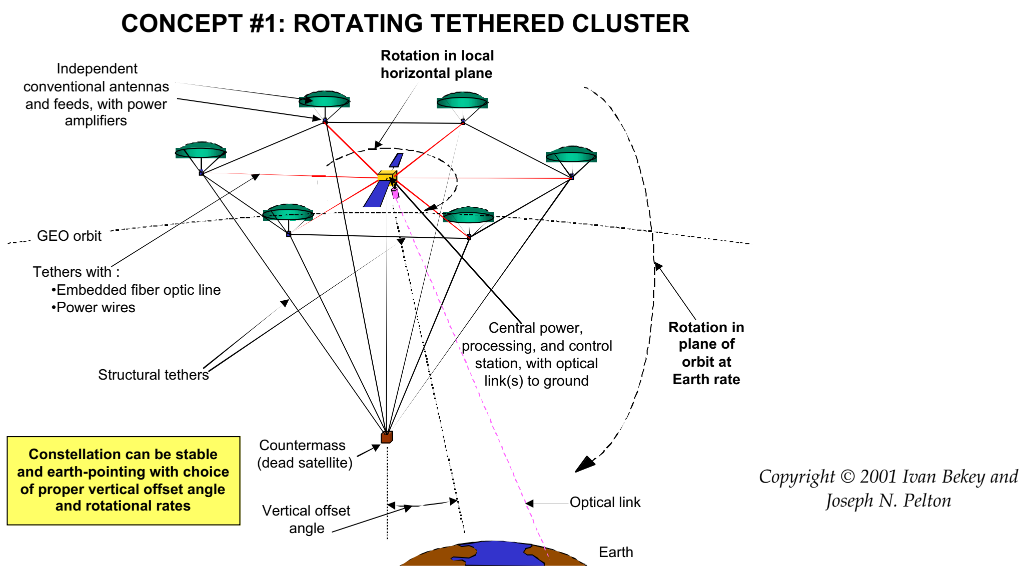

FIGURE 4

Hyper Satellite Concept 2

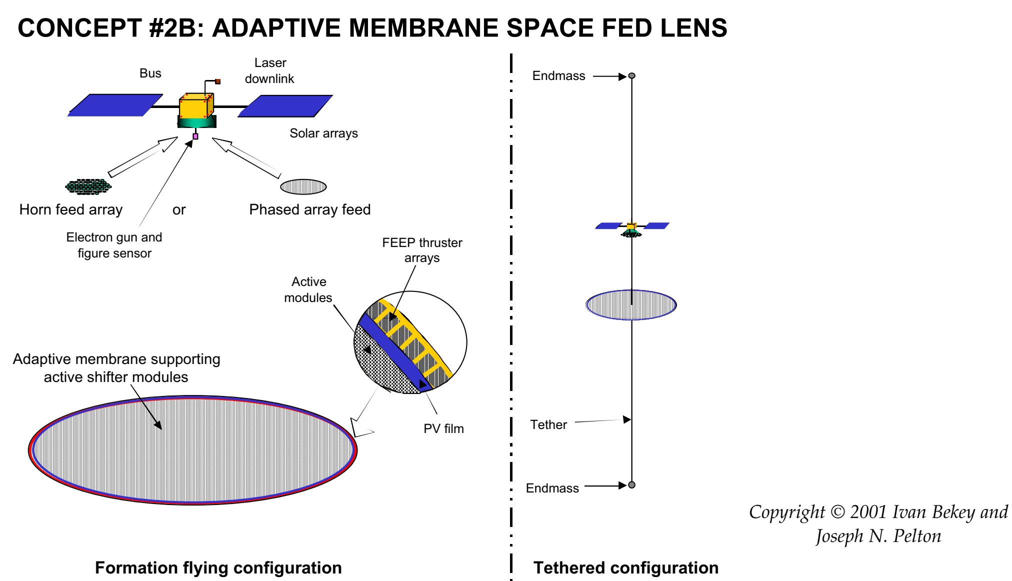

FIGURE 5

Hyper Satellite Concept 3

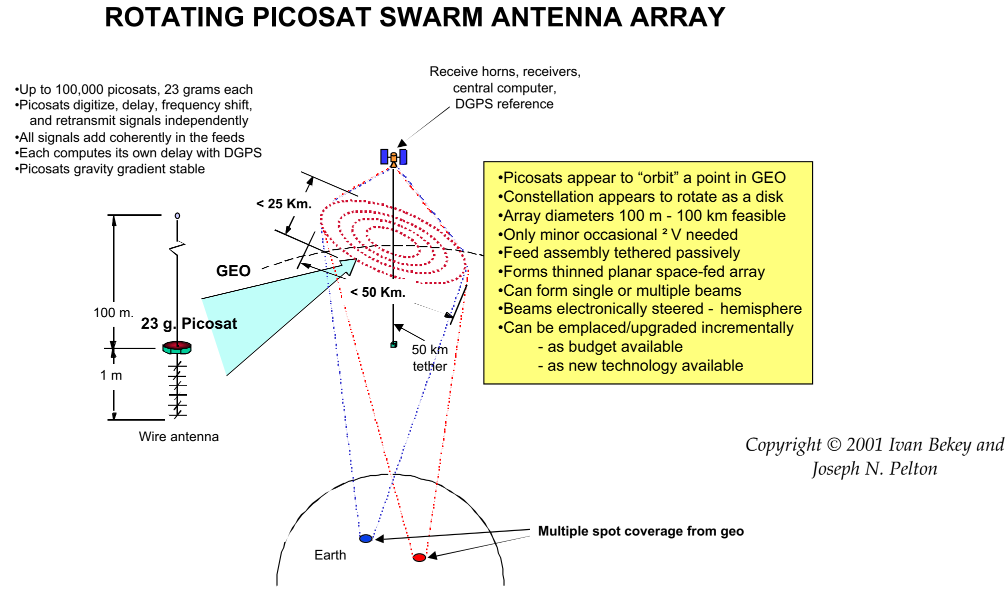

provide broadband services to cities. Such systems would be capable of very high system capacities with high cost-efficiency and with a highly desirable feature of low transmission latency. This concept is shown in *Figure 6*.

Only a network of a dozen or so optical satellites equipped with intersatellite links as well as optical links to a system of hundreds of HAPSs could provide mobile and broadcast services to create a global information network capability. If the HAPSs could be deployed at altitudes as high as 32 kilometers (or 20 miles), then they would be well above the atmospheric disturbance and rain attenuation that makes space-to-Earth communications difficult.

Further, the HAPSs would be equipped with enough on-board power to operate at very broadband, extremely high frequencies (EHFs) effectively. This is particularly true since the look angle would be highly favorable and the path-loss considerations minimal for the high gain, multi-beam phased array antennas that could be deployed on-board

FIGURE 6

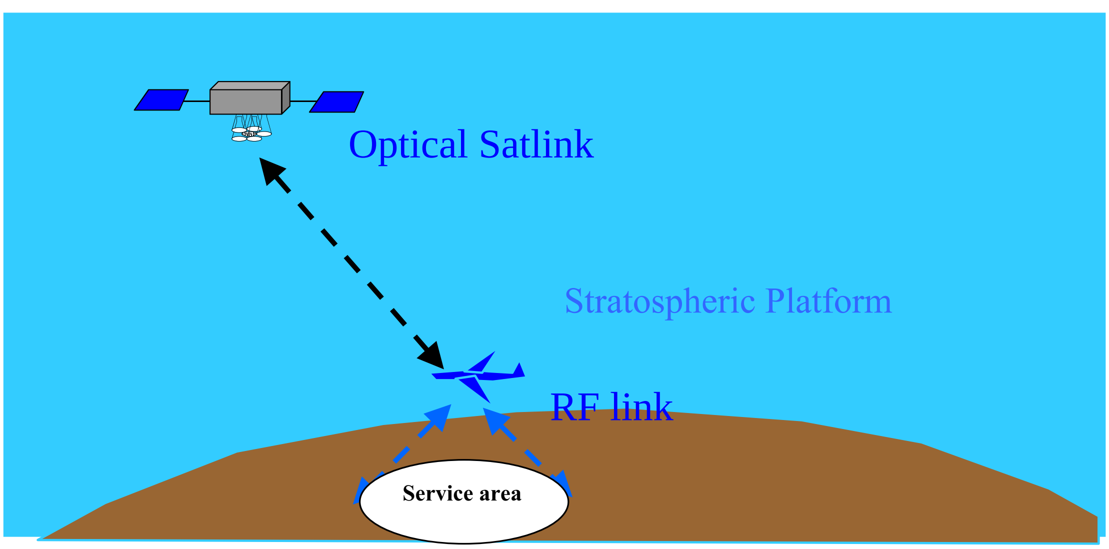

HAPSs by 2015. The detailed engineering of such a system is beyond the scope of this study, but a preliminary review of such a hybrid system is presented in *Figure 7*.

The overall consideration is that while the hybrid satellite/HAPS network would likely be cost-effective and address the latency problem for many applications and services, it is not certain that such a network would have sufficient capacity to address the rapidly growing demand for broadband mobile services. In short, such a hybrid system might be needed in addition to terrestrial mobile wireless systems in the 60 GHz and above bands as well as geosynchronous–Earth orbit (GEO) platforms.

In all of the aforementioned advanced architectural concepts, very broadband communications, including broadband mobile communications to portable (and even wearable) micro-terminals, would be technically (and likely economically) feasible.

Summary and Conclusions

The field of satellite communications has grown tremendously in terms of new technology and service capabilities over a span of nearly 35 years. The increase in capability from Intelsat 1 to Intelsat 9, for instance, represents a gain of well over a thousand-fold, but for satellites to maintain their economic and technical relevance in the 21st century, dramatic new gains are needed. The total new architectures discussed in this article show that "technical breakthroughs" are possible to create new hyper satellite systems a decade into the future. These system could produce not only major increases in throughput, but also breakthroughs in cost-efficiency, so that 2 gigabits of information could be delivered at a transmission cost of about $0.01.[1]

Selected Bibliography

Bekey, I. and Mayer, H., "1980–2000: Raising Our Sights for Advanced Space Systems," *Astronautics and Aeronautics* (April 1977).

Bond F., "Final Briefing of GEOPLATform Feasibility Study," NASA Contract NAS 8-32881, The Aerospace Corp. (June 18, 1979).

Bravman, J., Fedak, J., and Giffin, G., "Infrastructure on Demand—Global Fiber-Optic Capacity from EHF Satellite Communications Systems," Proceedings of the Fifth International Workshop on Satellite Communications, SS/GII (Ottawa, Canada: June 15, 1999).

Cuccia, C.L., Lovell, R.R., and Edelson, B. I., "The Evolution of the Geostationary Platform Concept," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-5, No. 4, (May 1987), 601–614.

Edelson, B. and Morgan, W. "Orbital Antenna Farm," *Astronautics and Aeronautics* (September 1977), 20–27.

Iida, T. and Suzuki, Y., "Communications Satellite Research for the Next 30 Years," Proceedings of the 19th ICSSC Conference, Paper No. 2001-233 (Toulouse, France: AIAA, April 17–20, 2001).

Landovskis, J. and Rayment, S., "Emerging Broadband Satellite Network Architecture and the Major New Developments," Proceedings of the Fifth International Workshop on Satellite Communications, SS/GII (Ottawa, Canada: June 15, 1999).

Pelton, J.N., "The Future of Telecommunications", *Scientific American* (April 1998), 80–85.

Pescovitz, David, "Wired for Speed: As Chips Shrink, Researchers Look to Optical and Radio-Frequency Interconnects," *Scientific American* (May 2000), 41–42.

Stuart, Janet Gleave, "Distributed Parameter Control of a Large, Flexible, Stiff-Membrane Phased Array Using Thin Film Piezo-Electric Actuator Patches" (Ph.D. Thesis, Dept. of Aerospace Engineering, University of Colorado, 1994).

FIGURE 7

KEY OBSERVATIONS ON A HYBRID HAPS/GEO SATELLITE NETWORK

- Such a hybrid system largely overcomes problems of latency for regional and intra-city communications, but interregional and international communications would still experience some 250 milliseconds of delay.
- HAPs would function as optical/RF transformers. (Thus, an optical signal can be sent to and from a geosynchronous satellite from a stratospheric platform.)
- HAPSs, in comparison to geosynchronous satellites, substantially reduce the problem of rain fade in Ka-band and above because of high power, favorable look angle, and reduced path loss (e.g., 35 kilometers maximum versus 35,800 kilometer altitude).
- Optical communications satellites can be deployed in multiple locations for space diversity, and thus direct space-to-Earth and Earth-to-space links would be available as an emergency "fail over" system.
- High levels of frequency reuse can be achieved by HAPSs (with well over 1,000 beams and up to 300-fold frequency reuse) for such applications as broadband mobile communications. Simpler systems could be deployed for entertainment and broadcasting applications
- HAPSs could serve localized needs, and optical communications satellites could serve inter-regional needs.
- High-altitude aircraft have now been developed and prototypes are in operational testing, and thus represent a higher state of R&D and feasibility testing than most geoplatform concepts.

Supler, Curt, "Nature's Way Might Be Path to Smaller Computer Chips," *Washington Post*, June 8, 2000, A-2.

Takats, P. and Garland, P., "Evolution of Regional Multimedia Satellite Architectures for Future Internet Services," Proceedings of the Fifth International Workshop on Satellite Communications, SS/GII (Ottawa, Canada: June 15, 1999).

Books and Reports

A Collection of the 18th AIAA International Communications Satellite Systems Conference Technical Papers (Oakland, Calif.: AIAA, 2000).

A Collection of the 19th AIAA International Communications Satellite Systems Conference Technical Papers (Toulouse, France: AIAA, 2001).
An International Forum on Space Solar Power: Conference Proceedings (Washington, D.C.: NASA, January 12, 2001).

Richard T. Gedney, Ronald Schertler, and Frank Gargione, *The Advanced Communications Technology Satellite* (Mendham, N.J.: SciTech, 2000).

National Science and Technology Council, *Networked Computing for the 21st Century* (Washington, D.C.: White House, 1999).

Joseph N. Pelton et al., *Global Communications Satellite Technology and Systems* (Baltimore: Md.: World Technology Evaluation Center, 1999).

Joseph N. Pelton, *Satellite 2001: The Transition to Mass-Consumer Markets, Technologies, and Systems* (Chicago, Ill.: International Engineering Consortium, 2001).

Proceedings of the Sixth Ka-Band Utilization Conference (Cleveland, Ohio: NASA Glenn Center, 2000).

Notes

1. Please note that these would be only the "cost" and not the "price" figure for satellite transmissions. Further, these cost figures do not include user terminals, ground switching, or value-added services to end users.

3G in Latin America: Building the Network

Willy Schotte
Senior Director
Bechtel Telecommunication

Third-generation (3G) wireless systems promise advanced services and revenues for the operators. Nevertheless, questions remain. Is your customer base and market ready for the new services? What 3G standard is best for you and your market? Should you wait for Internet protocol (IP)–based systems? Is an interim technology such as general packet radio service (GPRS) a good solution for now? Whatever your analysis, the switch from second generation (2G) to 3G is a major technological challenge; you will soon need major changes to your systems and organization.

Wireless companies in Latin America, including countries such as Brazil, began the cellular race late. They started late and faced the challenge of constructing networks and systems in less than half the time of typical U.S. companies. However, they have performed well and continue to build out and introduce new technologies. Cellular phones are becoming essential to conducting business and are very much a part of life. For the most part, carriers have built their networks and now are changing the business focus from providing coverage to providing capacity and features as quickly as possible. Now, the focus is on marketing and the introduction of new services and applications. The challenge is how to get and keep customers while maintaining profitability.

Today, the world is talking about 3G (third-generation mobile systems), and operators cannot sit back and watch their competitors take the lead. Operators need to say READY, SET, GO—and begin the race to implement the new 3G networks. The world is ready for high-speed data, music, and video on portable devices. Some believe we will still talk on portable phones in the future. The hype about 3G reminds me of the hype about integrated service digital networks (ISDNs). I really hope that 3G exceeds the market success of ISDN. Today, around the world, carriers are taking a sober look at the 3G revenue potential and are carefully analyzing the cost to implement it.

What Is the Best Approach for Carriers in Latin America?

One can correctly conclude that every flavor of technology, frequency, and service provided is present in every country, with very little roaming possible with advanced services. Can an operator remain with short message service (SMS) only, or is a combination voice and cellular digital packet data (CDPD) necessary? Several options in a dual-mode/dual-band handset are available. Eventually, if we wish to have full transparent roaming, we must align with the most common technology and with the most universal and open standard. That may be some version of 3G. Today, many carriers are taking the interim approach to network evolution and are moving towards 2.5G and GPRS. In several countries, the operators concluded that data rates around 40K are sufficient to meet the demand of the market for the next few years. They plan the move to GPRS and will use GPRS as a training ground for the new data world to which their operations have to adjust. The move to 3G may be less painful with 2.5G operation experience. However, the capital cost for 2.5G may not be recovered before the next change-out to 3G is required!

READY, SET, GO—Building the Network

While considering the move to the next generation, carriers are underestimating the efforts involved in moving to the new technology. The following are some points that could help operators avoid mistakes when building the next step in their network.

The variety of vendors, frequencies, and services does not allow me to be very specific on the issues; hopefully, there is sufficient information to trigger "opportunities" in your system. The new technologies are universal in the way in which they can be implemented. Any installed digital service can migrate to 3G. Any frequency band can be used for 3G. What a system!

This implementation could be easy—if one has a new license and if ample bandwidth is available. It is not so easy if the new service has to be implemented in the current operating system and has to reside within the current allocated spectrum. If major changes to the frequency plan are required, and spectrum clearing is necessary, engineers can expect serious challenges. The complexity and cost of spectrum clearing is directly related to the bandwidth to be cleared. Additional construction of cell sites may be required to protect the service during the installation phase of the new service. If the regulatory agencies have allocated new spectrum where current microwave systems are operational, spectrum clearing may add additional cost to the implementation, especially if the new user is forced to relocate the current systems.

Many alternatives can be considered if additional spectrum is assigned. One alternative is pre-seeding of the market with dual-mode/dual-band phones. Be warned, several operators have reported that in some areas, even though the voice quality is acceptable, the data-transfer quality is not. This is not the end of the challenge. The installation of additional equipment in cell sites requires additional floor space. Those upgrading to GPRS should consider looking at the requirements for migrating through to 3G. The GPRS implementations for some vendors are simple and do not require major equipment changes; however, major additions may be required for the next step to 3G. Developing current, correct site documentation will save you money in the long run.

The additional demands for alternating current (AC)/direct current (DC) power and the increase in air-conditioning requirements add to the complexity. On most rooftop sites, space is limited, and, in some cases, expansion of the systems has maximized the rooftop loading. This may require a review of the structural parameters. Grounding also needs to be reviewed. Cell-site equipment most likely runs on DC power and is well screened from outside influence by the AC/DC power systems. Now several vendors are introducing direct-powered AC equipment sites, such as data routers and hubs located next to the DC equipment, and are introducing equipment that is not typical. Several violations of grounding rules have been observed that expose the site to possible damage by power fluctuations and lightning strikes.

Site acquisition becomes very difficult in cities. Most rooftop prime locations are taken, and if not, we find a radio frequency (RF) nightmare on the rooftops. Safety for personnel becomes an issue, and interference is present. When a tower is feasible, co-location may be the only option; here, structural loading is again of concern. Additionally, the rooftop and tower owners know the value of the roof or tower!

If site acquisition is successful, we then have to deal with increased complexities in the zoning rules and permits required for construction and changes to structures. Moratoriums on towers and antennas are present in several jurisdictions. Even a minor change in the dimensions or visual change of an existing antenna requires weeks of waiting for zoning approval. The co-location of antennas on rooftops or towers becomes a major issue, considering the additions of new frequencies and transmission of high-density data signals. Inter-modulation to your current systems and to the other co-locators should be analyzed.

New equipment such as the diplexer is introduced to the RF path of working systems. This will change the losses in the RF path and the characteristics of the site. Coverage and quality may be affected. The introduction of wideband signals with high data rates requires a complete review of the RF path. Several devices in the current systems—such as filters, combiners, lightning protection, and test systems—may have bandwidth restrictions, and change-out would be required. Several of these devices cannot meet the specifications required for the high-speed data.

If the 3G signals are added to the existing RF path, the power specification of the equipment in the chain has to be reviewed. In some cases, equipment from different vendors may be used. The interaction between this equipment must be carefully analyzed.

Today's transport networks are mostly built on channel-based T1 or E1. The system tolerates minor impairments without major impacts to the service. Compression, drop, and insert are used frequently. Most existing wiring is below standard for the new applications. The microwave spectrum is open, and the local carrier can meet the T1 or E1 demand. The move to wideband signals and high-speed data demands major changes to the network. T1 or E1 cannot be channelized. Several products cannot meet this requirement. Drop and insert is not possible. Wiring must be upgraded considerably; this may be a challenge to the service provider. If the system depends on a microwave system, a review of the specifications for non-channelized high-speed data is recommended.

Remember that under the 2G technologies, most mobiles are dormant and observe the control channels for activity. They use minimal bandwidth while in the idle mode. The 3G equipment could be very active receiving data while at rest on your desk and be perceived to be idle.

In a system with high data density, the quality of the network becomes important. Poor bit-error rate (BER) performance and the resulting retransmission of data packets may decrease the useful capacity of the system and bring unacceptable quality of service (QoS). The measurement of "lost packets" is also essential. The traffic-engineering rules being tested for the new systems are complex. To guarantee QoS, it might be that meeting the objectives on voice quality and dropped calls may not be sufficient indicators for good service. Can you see the impact of missing packets in critical applications?

How Can We Minimize the Opportunities for Problems?

The following are of equal importance:

- Clear market and marketing objectives
- A good entry plan
- Pre-construction surveys and site inventory
- Review of structural loading
- Well-developed RF plan
- Well-developed network plans (for voice and for data)
- Program management and logistics for large builds
- Integration and cutover plans
- Operations support system (OSS) adapted for the voice/data systems

For those of you that plan to build systems and then later consider upgrading to 3G, I recommend that you manage this move as a new system build. Most of the problems that you solved in your first build will return, and then some. The need for experience is key. Training your staff on the new specifications and system requirements is essential. Training should also focus on data systems. Technicians at your cell sites will now need to be skilled in information technology (IT). Carefully select your OSS. Don't forget you manage data—not channels, circuits, or voice.

The change from 2G to 3G is a major technology change. It requires major adjustments of your people and support systems. It is like a new language—you have to learn it!

My final comment is that you should not forget to watch the market and be prepared to make prompt adjustments in your course. Building a network is a big challenge. We are building the network to offer services to our clients and to anticipate their needs and their willingness to pay for the service. The latter may be the biggest challenge of all.

Mapping Mobile Commerce to Network Support Systems Requirements: An Overview

Nicholas C. Stanley
Principal
Lecoma International, Ltd.

Summary

The mobile telecommunications industry is now making a vast bet on the future: that there exist compelling reasons for customers to use next-generation mobile networks in ways that will produce revenue streams well above those for voice. While revenue in and of itself is good, the carriers also look to next-generation services to decrease churn by making users more dependent on the carrier's services above and beyond generic voice. Reducing churn and pushing more content through the system's infrastructure are both vital to the bottom line.

While early indications for 2.5G and 3G are mostly a product of speculation in the trade press, less attention has been directed to the task of "monetizing" customer behavior—that is, the means to translate consumer use of services to cash flow.

This paper[1] will attempt to model various consumer behavior and identify what business support system (BSS) needs could be deduced from the payment models. There appears to be a strong business case for continued use of third parties in the payment and settlement process, much as there are in conventional retail services. Indeed, mobile commerce may show many parallel elements to conventional commercial structures, which leads to the conclusion that the dominant industry players in consumer retail payment collection and settlement may have significant influence upon the fate of the wireless industry—a new force has entered the arena.

Contemporary Projections and Conventional Wisdom

As this paper is written in the winter of 2001–2002, both the popular and trade press are frequently populated by articles on next-generation services: Some are optimistic to the point of hype, while others forecast financial meltdown.

The future is unknown. Certainly the most optimistic models are based on wishful thinking but do have a basis in fact: Mobile services have historically vastly exceeded expectations (in the early days of cellular radio, the total number of cellular phones in the United States was expected to be less than one million). The cynics can point to massive failures of business models still shaking the industry (Excite@home being one).

Nonetheless, a few facts put things in perspective.

- *Consumers find ways to do new things.* Witness short message service (SMS)—originally almost an afterthought in the Global System for Mobile Communications (GSM) model. Estimates are that more than 200 billion SMS messages will be sent this year[2]. If the carriers capture just U.S.$0.01 (or even €0,01) for each message, that equals U.S.$2 billion or just under that in euros. Quite a nice bit of cash.
- *Things unique to mobiles become lifestyle component*. Ring tones, which can use either existing SMS–type technology or new 2.5/3G systems, are expected to bring in about U.S.$300 million in Europe in 2002[3] and at about U.S.$1.00/€1,00 each—that's a lot of beeps and burps.
- *New services start in one segment and move outward*. One doesn't have to repeat here the success story of i-Mode in Japan, where DoCoMo is reaping the benefits of the attraction of mobile content to demographics far beyond the original teenage early adopters.
- *The fixed Internet provides a guide, but perhaps not a roadmap.* Mobile content as a whole (which includes ring tones, multimedia alerts, electronic greeting cards, etc.) is predicted to catch on as SMS has done and, using projections from the Internet use of "push" and "pull" technologies, may reach significant levels as next-generation services are rolled out.

Many, many other examples can be pulled from the media. The purpose here is not to establish the credibility of the argument that new services will mean new revenue streams—that much is given. What is not so clear, however, is how much of that cash flow is going to be applied to the amortization of the capital investment in the infrastructure that makes the revenue possible. It's not beyond the realm of possibility that the carriers could find that their BSSs are unable to cope with rapidly emerging services. Those services, by definition, are in high demand by consumers, and somebody will find a way to deliver them. If the carrier's BSS cannot capture some part of the value for particular transactions, the remaining cash flows from other services must pick up the load on the marginal costs involved.

The old paradigm was "If you can't bill it, kill it." The implication was that the carrier had a choice in the matter. No longer. The carrier may not have that luxury.

End-User Models

The need for carrier's BSSs to evolve and develop in concert with the next generation of services is hardly news. Industry experts have been producing analyses and proposing solutions for some time. One thing that is given scant attention in these articles is the actual segmentation of the consumer market. This is natural, as most experts come from the billing-systems industry, where the "consumer" is actually a "subscriber" and as such is a fairly homogeneous class. All cellular subscribers act pretty much in the same manner: They talk;, they roam; they use a few vertical services. Up to this point, most billing companies were concerned with how to gather information on events, rate those events according to a set of rules, and then produce a bill for the services and perhaps a settlement report for outside partners.

Now, the challenge for the carriers is to find a means to evolve their billing systems in such a way as to capture as much revenue as possible from the flow of next-generation services and content delivery. Sophisticated and complex models for partnership settlement, revenue sharing, and so forth are spun out. Complex software tools are being developed and intensively marketed.

While certainly far from trivial, such a picture misses the point. Consumers are becoming heterogeneous, and no matter how comprehensive the BSS becomes, it is fated to be only one component of several. The diversification of the consumer base brings with it a strong need for the carrier to support the personalization of service offerings. Every customer will have a slightly different selection to fit his/her preferences. While the aggregate choices may be fairly uniform with a customer segment, carriers must take it as a given that no "one-size-fits-all" billing model can exist.

We postulate here that in the next generation of mobile services, consumers will decide for themselves how they wish to pay. Just as they do now in a store, they will demand flexibility, and the degree will depend on the value of the transaction relative to the relationship that the consumer has with the carrier. We have, somewhat simplistically, divided consumers into four segments:

- *Lifestyle Consumers:* The prototype is the Gen-X young consumer using a mobile as an integral part of a fashionable lifestyle. Voice, images, ring tones, "texting"—whatever is new and innovative, they try it. As new ways of communicating are found, the carriers and service providers may find themselves being led, not leading. How do they pay for all of this? Whatever it is, it has to be simple. And affordable.
- *Convenience Consumers:* This group uses mobile services that provide a function or a feature to them that just makes life easier. The soccer mom getting directions to a new school for a meet, the business traveler checking for flight information in the taxi, maybe getting the occasional alert from a stockbroker. The dominant theme is using the unique capability of mobile access to information, but on a demand basis. Spend a bit more? Sure, it's worth it.
- *Deliberate Consumers:* A more intense blend of the lifestyle and convenience consumer segment. The deliberate consumer has built mobility into his or her profession. Examples abound. The road warrior with regular access to his e-mail and voice-mail via integrated messaging. The real-estate broker who is also an active day trader, so she requires real-time quotation and access to trading. The deliberate consumer makes conscious, rationale choices on what services he or she needs and how much to pay—and how to take the money from their pockets.
- *Enterprise Consumers:* This class has received remarkably little attention in the trade press regarding their impact on 2.5G and 3G networks, perhaps because it's far less glamorous. But it's far more lucrative. An enterprise consumer is an individual that uses mobile services but has no personal connection to the payment process, although they may be aware of the costs. Examples here are employees using mobile services under a company account for voice, access to enterprise information (either via a mobile portal or the Internet), and specialized information services. An enterprise customer can also be a deliberate, or even convenience, customer. For example, getting road directions while on a sales call. Personal use outside of business is almost a given and would be hard to prevent, and thus the implications are many.

Consumer Payment Model: Money Makes the Map

We now come to the exercise of mapping the consumers onto payment models. Note that we have not mentioned the type of network used—2.5G, 3G, i-Mode, wireless local-area network (WLAN). The underlying network infrastructure is, at this point, almost irrelevant. The driving choices here are as follows: Is the consumer aware of the cost of the services, and, if so, what are their likely choices for settling their bills? And, given those choices, which ones are likely to dominate? Given these choices, the implications for the BSSs within the carriers are starting to take shape.

Likely Choices

Figure 1 shows a projected range of payment models. Note that we have not assigned any particular customer segment to the choices. The figure is presented in the form of a decision tree, with the branches determined by what choices the customer would have to make. The endpoints of each branch are the possible payment models. Each model has certain implications for the carrier BSS, and some overlap as indicated.

The payment choices defined in *Figure 1* are somewhat arbitrary for the sake of simplicity. While it is ultimately true that each and every transaction must find its way to the end of a decision tree, we'll stick to the most obvious choices.

- Enterprise charging is a direct billing of services from a carrier to an enterprise, which in turn allows members to use the carrier's services under some type of contractual, negotiated term. An enterprise can be a company, a nonprofit, a government entity, etc.
- The Inclusive/Package payment branch is similar to existing "bucket" plans in which consumers can buy voice minutes. In this case, perhaps a bucket of voice minutes comes with 200 SMS messages and five ring tones a month. Anything over that goes down the next path.
- The Carrier Billed/IN Wallet path is for transaction payment where the carrier's own systems handle the financial settlements, either by adding a line item to a recurring bill for a subscription or by decrementing a store of value within the carrier's network. Prepaid services fall into this category if the prepaid services can be charged at a discretionary level by the consumer specifically for transactions, not just voice.
- If the carrier has arranged with a third-party clearing or financial-settlement house to charge consumers for the services via an external mechanism, typically a credit card in the United States or bank debit in Europe, then the transaction takes the next branch. Here, we differentiate between a third-party system in which each and every transaction is billed and one in which charges are aggregated then billed to the external payment system at some period (every month) or at some value (every $25).

Figure 1 can be read as follows, reading left to right, top to bottom.

- *Branch #1:* If a consumer is not aware of the payment/settlement for the services, then either those services are included as part of a package (e.g., 100 free SMS a month, hard limited; access to a carrier's location service is part of a premium package) or the consumer is an enterprise customer, and the bill is settled by their employer.
- *Branch #2:* Here the consumer does have a choice. The two fundamental choices are 1) to have the carrier do the billing via a subscription (charges appear on the monthly bill) or via some type of stored value in the network (this would apply also to prepaid), or 2) to have a third-party payment system.
- *Branch #2, Sub-Branch.* If a third party is involved, they may be either for single, one-off transactions of presumably higher value or may be an aggregator that collects a number of smaller monetary charges and then takes payment from the consumer at some future time based on a commercial scheme (monthly or when the collected charges hit a set amount.).

The key point for *Figure 1* is simple: *Every consumer/customer for a carrier will follow a decision tree of similar format.* The decision points will be determined by their personal preferences and may change over time. Also, the flow down the decision tree may be different for different individual transactions. For example, a subscriber may have text SMS subscribed but pays via a consolidated third-party for stock quotes. However, it's likely that a given segment will show a strong affinity for certain choices. *Therefore, the implications on BSSs can be driven from consumer-segment behavior.*

FIGURE 1

Payment Models for the Consumer

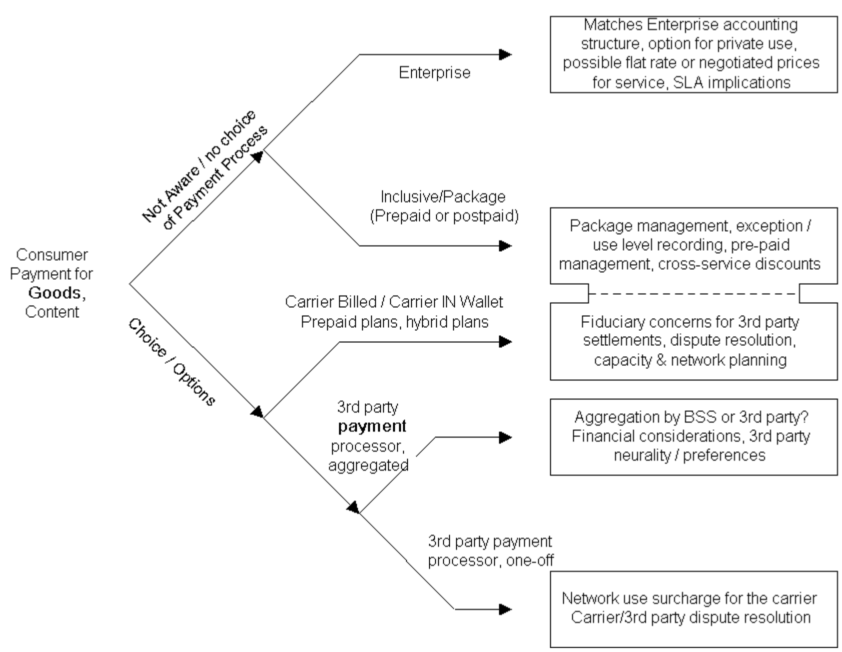

The implications on the carrier's BSS are shown in the boxes on the right in *Figure 1*. These are hardly inclusive, and each box could be deconstructed in a complete analysis far beyond the limits of this paper. The concept, however, remains. The cash flows from each endpoint are distributed amongst multiple entities. The sum of the cash flows for the carrier are what will drive the carrier's revenue for content and services and thus are the only source of amortization for the infrastructure investments.

Limitations and Boundaries

Implicit in the model is that some limits will exist on consumer behavior. These limits can be soft, in the sense that consumer behavior for wireless services will match that for any other discretionary purchase. Other limits maybe hard, in the sense that a carrier may choose to set financial or monetary restrictions based on its own business models.

An example of some of these limits are (in no particular order) as follows:

- A carrier may choose to limit the amount of external services (e.g., directions, stock quotes, ring tones) from external suppliers to no more that $XX per month if those services are to be direct billed.
- External content providers that are paid using a subscriber's credit card may not accept transactions lower than a certain amount, demanding that the network or a third-party consolidation of disparate charges take place.
- Carriers may offer a stored value service ("wallet") but, for fiduciary or risk reasons, may limit the amount that may be deposited.
- Taxation for out-of-jurisdiction transactions may force certain events to a given payment model.
- Individuals will choose logically to set their own choices. This is shown in the next section.

There are many other factors. These may make the decision tree more complex, but in the end, all consumers must travel down one route or the other.

Mapping Models to Use

Figure 2 is an attempt to map the various simplistic models in *Figure 1* onto revenue. It's difficult to predict the amount of money associated with these future services. Instead, we have used a normalization process.

The horizontal axis in *Figure 2* relates the value of each individual transaction to the basic wireless service charge. The rationale is that the consumer will treat an incidental or casual use of a service differently depending upon its relative size to their actual recurring or baseline charge. For example, a fifty-cent ring tone bought once in a great while is readily accepted to be billed on their U.S.$50.00 monthly subscription under "miscellaneous" (normalized value = 0.01), while the booking a payment of a U.S.$75.00 train ticket (normalized value = 1.5) is something else. The normalization takes into account that high-value customers (say, U.S.$500.00 a month) would have a different view than the youth customer trying to keep things at less than U.S.$40.00 a month because that is all his father allows.

The vertical axis in *Figure 2* is again normalized, this time taking the ratio of the total monthly consumption of mobile-commerce services, whatever the payment method, to the baseline subscription. A ratio of 1.0 indicates that the total amount of discretionary content and services is equal to the baseline subscription charge. The period can be the billing period (for post-paid) or the consumer's own budget period

FIGURE 2

Payment Models As a Function of Monetary Value Related to Wireless Service Value

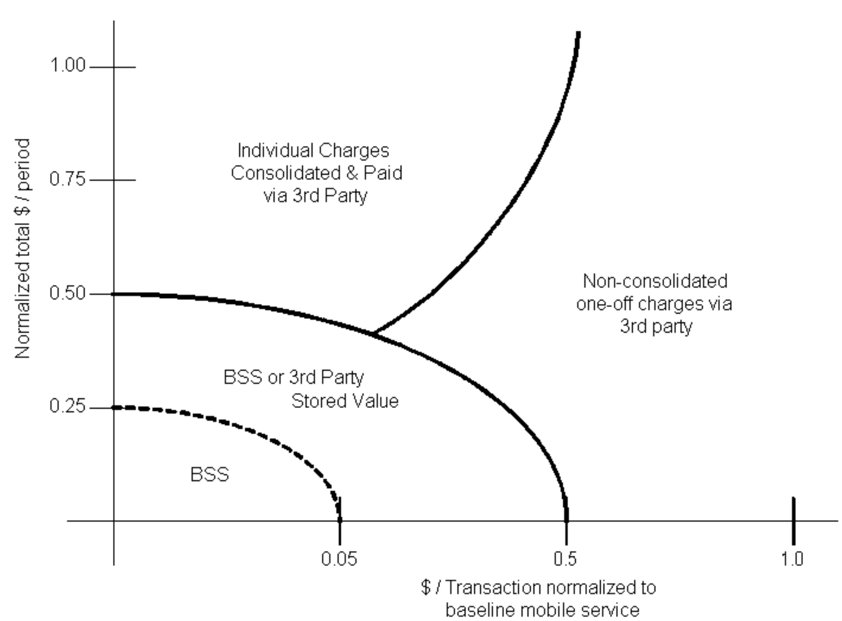

for prepaid, however informal. The rationale is parallel to the first. Namely, a customer that has mobile commerce transactions, however many, that are still a small part of the bill will prefer certain payment methods to another with the same set of transactions but instead are very significant compared to the "normal" bill. We have arbitrarily taken some breakpoints at 25 percent and 50 percent for illustrative purposes.

Figure 2 presumes that there are certain boundaries to types of payment models.

- "BSS" refers to the routine billing or charging of things within the core BSS of the carrier, post-paid or prepaid. The boundaries are dotted, as this area is probably the least defined.
- As the discretionary use of mobile-commerce services increases, the carrier may choose to enforce a wallet or stored value to control post-paid risk and for handling prepaid subscribers who may wish to segregate payment for discretionary purchases from basic communications. The intercept points on the axis are highly debatable, but the basic shape is probably similar—the implication is that the internal BSSs of the carrier will be handling only a subset of all possible transactions. The carriers are also keenly aware that there are fiduciary implications for them as the amount of stored value increases. For nationwide carriers, even a few tens of dollars in millions of accounts represents a significant figure. Also, they must take care how those funds are used and stay away from any services that begin to look like banking (such as transfer of funds to another person or entity independent of a purchase transaction).
- Outside the area of the BSSs, we see two distinct areas. One, in the upper portion, bounded by the y-axis, is the area where third parties consolidate multiple smaller transactions into reasonable units. There may be many such providers, servicing different segments or different services—ring-tone portals, travel sites, telematic services, etc. The other area is where there is a significant monetary purchase that is conducted over the wireless infrastructure but has a single discrete payment and settlement process. The obvious is an e-commerce purchase over the Internet that just happens to use a wireless connection: buying an airline ticket, a book, a CD, etc.

Payment Models and Customer Segments

The last step is to map the customer segments discussed earlier onto the payment map. This is shown in *Figure 3*. It is obviously very much a subjective process, given that little data is available and that what is known is a carefully guarded secret of the carriers.

There are two key points in *Figure 3*. First, each consumer segment must fall into at least one area of the payment space. Thus, the total customer base can be mapped onto one or more payment mechanisms, and this mapping is determined by consumer behavior. Second, the relative value to the carrier will be determined by the marginal revenue obtained from each customer as well as the marginal cost to produce that revenue.

The second point is an interesting one to consider. A model which attempts to push the boundaries of the BSS outwards, on the theory that the more revenue captured the better, runs up against consumer behavior as well as the increasing cost to produce that revenue. Expansion of a carriers' BSS to encompass aggregation and collection may

FIGURE 3

Consumer Segments Mapped on Payment Models

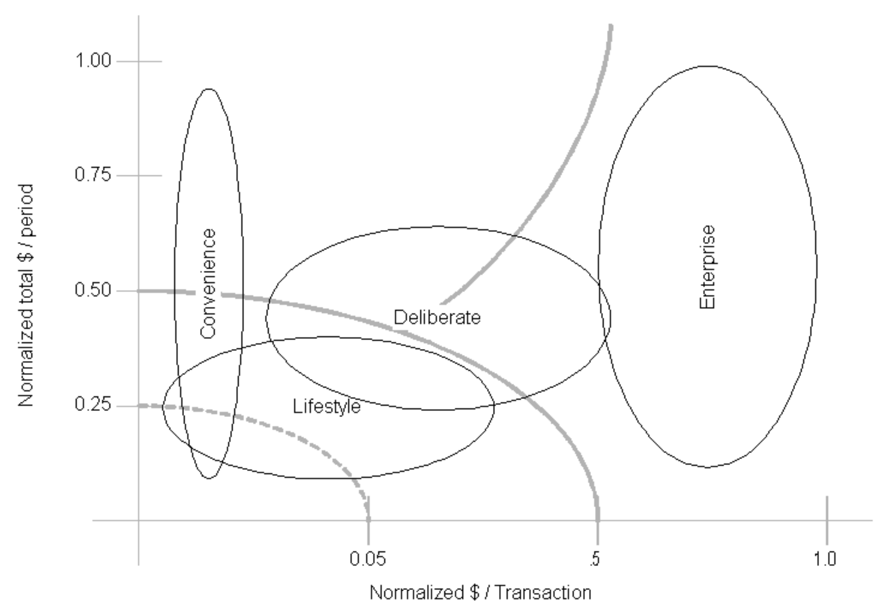

be best left to those who are best at doing such things. The key is how much of that segment is left in the BSS and what falls elsewhere.

Conclusions: Implications for BSS in Carriers

The impact on the carrier's BSSs can only be covered at a very high level, given the level of detail in this paper. *Figure 1* presented a very concise set, which may be readily expanded from the basic ideas presented. Clearly, a payment method involving an enterprise customer must match the internal accounting structure of the customer. An intelligent network (IN) wallet of stored value within the carrier network creates a fiduciary responsibility on the carrier to properly manage and secure that value. Use of a third party is essentially an outsourcing operation with some of the characteristics of factoring. It has no direct counterpart to the retail environment, where every transaction stands alone, but it is also not a subscription relationship between the consumer and the carrier, either.

The primary difference in this analysis versus many in the media is that it has attempted to link customer segments to operations support system (OSS)/BSS implications and, in doing so, provides a rough model on how such an analysis could be undertaken with more rigor. Certainly it is not the only model and may, indeed, be only a means to provide a reality-check on other planning and business models.

Fundamental aspects relating to BSSs relate to consumer behavior, in particular behavior that relates to a changing view of mobile services. For the lifestyle consumer, mobile services are a locus of communications and exciting, different ways to interact socially. For the deliberate consumer, services become an enabling part of their professional lives. For the enterprise consumer, it becomes one more tool on the job.

Individuals charged with BSS and OSS planning have to consider the needs of multiple "customers" for their complex systems: the external customer (the one who pays the bills!), of course, but also internal customers (network planning, marketing, finance, sales) and partners (content providers, third-party payment houses).

The planners will have to wrestle with questions such as the following:

- *Granularity versus Cost:* At what point must an individual charge be maintained within the BSS? Example: A what point does internal aggregation of cheap services become mandated? Does the BSS have to track millions of transactions a month if they are worth only pennies each?
- *Customer-Relationship Management (CRM) Implications:* If a convenience consumer buys a movie ticket on their cell bill and later wants a refund, how is the settlement made?
- *Fiduciary:* If an IN wallet for payment of external partners resides in the network, what is the liability for record keeping/audit of the use of those funds? How much overhead does the accounting function take relative to the income gained?
- *Neutrality:* Will there be regulatory restrictions on the use of "in-house" aggregators by the carriers if doing so places third-party clearing companies at a disadvantage? For example, where the net effect is to favor the carrier's preferred suppliers who can clear payments via the in-house system while other content providers must pay a third-party handling fee. This may not be a concern for the lifestyle consumer, but you can be sure that the deliberate consumer will have something to say about it. The customer segmentation allows for a better handle on the pros and cons of each choice.

Predictions

Always a risky proposition, but we believe certain trends are taking shape that are consistent with our viewpoint.

Prediction #1: Big billing systems will stay, but will not longer be the only systems that touch the money. The carrier's BSSs and OSSs will have to migrate away from a single, monolithic system to a "best-of-breed" approach. The primary reason is simple flexibility—no one knows for certain what revenue model and consumer model will dominate. We suggest looking outside the telecommunications industry to world-class consumer-oriented providers, where best-of-breed CRM, enterprise resource planning (ERP), and financial systems are the dominant players.

Likewise, the process of aggregation, as discussed earlier, could be done by the billing system or by an external third party. Carriers must choose internal complexity over control.

Prediction #2: There's room for more players and for innovation in payment systems. There is a role for external third-party payment and settlement processors now and in the future. If for no other reason, carriers want to be carriers and not financial clearinghouses. They cannot match the sophistication of the credit-card networks nor do they have the capital reserves and transactional integrity of the banking system. Nevertheless, the interaction between the external parties and the carriers will have to be through the BSSs. We predict a tighter integration of BSS functions and external third-party payment systems. The ones that do the best job will see lower costs of implementation and a cost advantage in the marketplace. In a marketplace with multiple carriers and a vast array of services and content, this price/convenience advantage could be critical.

Prediction #3: Content will stay content. We present in the Appendix a very simple flow model for content. Over time, the economies of scale and the need to reach the particular segments of the mobile consumer market will re-establish some (if not all) of the distribution (middleman) elements shown. Content, in the widest sense, will migrate to a business model with the least amount of economic friction between the consumer and the source. While seemingly complex and inefficient, we posit that in fact the marketplace has shown them to be highly efficient; after all, in a capitalistic open economy, inefficiencies in distribution are a fruitful area for innovation. Therefore, whatever the model, the producers of content will be disconnected from the underlying carrier infrastructure and its BSS/OSS if they are not there already.

The entire area of the impact of next-generation mobile commerce on the carrier's infrastructure is extraordinarily complex, with enormous ramifications for the ultimate profitability of the carriers. Certainly they will survive, but at what level of profitability is unknown. Consumer behavior, and the understanding of it, is the place to start.

Appendix: A Content-Delivery Model

Figure A-1 represents an idealized picture of how content flows from talent to consumer. It's simplistic in that it may show functions that are not always present, nor does it take into account new business models unique to the Internet and mobile commerce.

Nevertheless, this model has existed in one way or the other for hundreds of years and survives because each element of the chain provides some type of defensible value or has been mandated by regulation/law.

Where new business models allow that value to be moved up (or down) the value chain, or bypasses regulatory-defined structures, that function can and does go away. In the special case of legal or governmental involvement, some business models may be restricted or declared illegal (e.g., Napster).

The model in *Figure A-1* can be read from left to right. Each content piece (in the leftmost box) relates to the downstream (toward the consumer) piece by reading across at the same level. Hence, it's Movies – Artist – Studio – etc. (with apologies to the independent film industry).

The diagram can be used to determine where the value chain relating to the delivery of mobile content becomes specific to the mobile environment or access network for the end consumer. For example, a ring tone may not be mobile specific until it is bundled by a media company. Granted, a ring tone is useless unless it's in a mobile handset, but as a item, it's just bits until it is on the market and offered to consumers.

Notes

1. The author wishes to acknowledge the contribution of Mr. Ron Faith of Qpass, Inc. (Seattle, Wash.; www.qpass.com), who provided ideas and guidance for this paper.
2. Brad Smith, "Proponents: There's Hope for SMS," *Wireless Week*, October 8, 2001.
3. Allyson Vaughn, "Mobile Music to Their Ears," *Wireless Week*, October 8, 2001.

FIGURE A-1

Content Value Chain

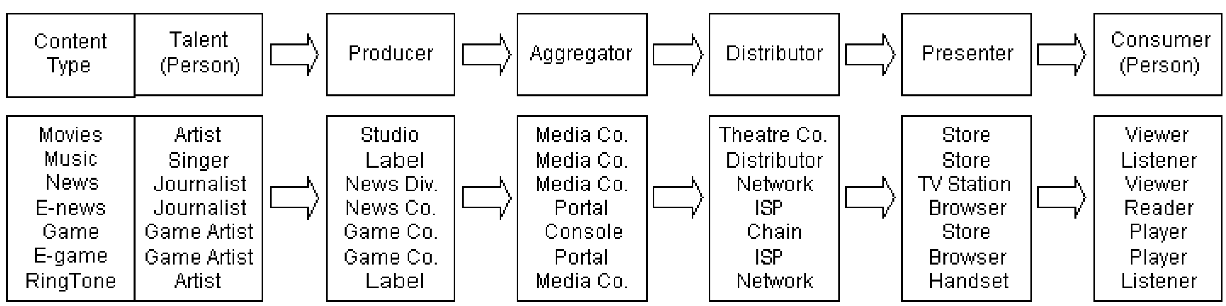

The UMTS Marathon: Maneuvers for the Long Haul

Patrick Tao
Vice President, Marketing
Tesaria, Inc.

The Global System for Mobile Communications (GSM), general packet radio service (GPRS), and Universal Mobile Telecommunications System (UMTS) equipment industry has changed. For brevity, I will use the term GSM equipment to mean equipment in the GSM, GPRS, and UMTS family. Where once an expanding market made room for many competitors to grow their business and sustain a break neck pace of research and development (R&D) investment in new products and new technologies, the economic downturn of 2001 and the contracting GSM equipment market is forcing all competitors to put spending under a microscope. This lesson is not limited to the GSM equipment market but applies to all telecom equipment markets. I will focus on the GSM equipment market to illustrate my thoughts.

The late 1990s were good years for the wireless equipment business. The industry grew 20 percent to 25 percent each year, with equipment revenues peaking in 2000 at roughly $56 billion. Long-range forecasts for the wireless equipment market still call for substantial growth once the economy recovers and wireless operators begin to invest in network expansion again. The picture in the next year or so is not so rosy.

GSM operators around the world have announced year-over-year reduction in capital expenditure (CAPEX) plans. From these announcements, many analysts have concluded that revenues in 2002 for GSM equipment suppliers will be flat or down by as much as 10 percent from revenues in 2001. In response, major GSM equipment suppliers (original equipment manufacturers [OEMs]) have been forced to shed product lines and cut staff to get their financial house in order.

The race between OEMs to bring new products to market has become a marathon where once it was an all out sprint to the finish line. Companies with the financial endurance to last through the downturn and into the eventual third-generation (3G) networks build-out will be the ultimate winners. Companies can no longer speculate on technology roulette. R&D budgets are being trimmed to support only those projects with solid, tangible, near-term, financial business cases.

Many observers expect a rationalization of the GSM equipment industry and the larger wireless equipment industry to which it belongs. The assumption is that at the end of this transition period, there will be fewer competitors for the GSM equipment market than today.

In this paper, I propose another business model to achieve a rational industry structure. The system specialist model is one that consolidates the horizontal supply chain so that current OEM suppliers can achieve the economies of scale necessary to compete; the industry achieves greater efficiency in the use of engineering and technical resources; and the industry sustains the product and service innovation needed to spur overall industry growth.

Wasted Product Development

The GSM equipment development business has been wasteful and inefficient. Each GSM equipment supplier today devotes hundreds of people and many millions of dollars to the development of GPRS and UMTS equipment. Most of these people are working on products with capabilities defined by the industry's standards. There are literally thousands of pages of standards specifications that have been produced during the past few years defining the functionality that every GSM equipment supplier must provide just to achieve the basic operation of the system. This mountain of technology specifications continues to grow faster than the ability of GSM equipment suppliers to actually implement those specifications.

In a thriving market, the duplicated efforts between these companies are buried under the overall positive results and are seldom questioned or dealt with. To stay competitive in a declining market, companies cannot afford this duplication. If the competitors in the GSM market had comparable revenue streams, the inefficiencies would affect all competitors evenly, but, in fact, the competitors in this market have widely different market-share and revenues. This gives the largest competitors a distinct cash-flow advantage for outspending and outlasting their smaller competitors.

Today, OEM companies are focused on cleaning up their balance sheets and generating immediate, positive cash flow. Gone are the days of long-term, "strategic" investments in products. The industry has returned to R&D spending guidelines such as "each product must be capable of generating revenues 10 times the development cost" and "at least $1 million in revenue per employee." Under these operating constraints, if broad competition among OEM companies is to remain, the smaller competitors must find new, cost-efficient ways to deliver products.

When we look at other industries with many sustaining competitors, we find examples of supply chains where product suppliers source their common technologies from companies that specialize in those technologies. In the GSM equipment supply chain, many companies have established successful businesses supplying technology components in this manner to the OEM suppliers; but we must go further.

Hardware platform suppliers have emerged in the last few years as the OEMs realized that they could use volume manufactured hardware to save time and money. There are now many choices of hardware suppliers for systems based on compact peripheral component interconnect (PCI) technology. On the software side, companies such as Wind River, Hughes Software Systems, and Trillium have provided common software protocols and functions to the OEMs that provide savings in software development.

The availability of technology components changed the way in which OEMs develop products. The OEM developer's job has become the integration of hardware and software from the component suppliers and the development of control logic to orchestrate the components into full working systems. These two jobs are where the duplication of effort exists today, and this is where the opportunity to fulfill the role of a system specialist supplier emerges.

When the OEM sources technologies from traditional components suppliers, the onus is on the OEM to develop the control logic for the components, to make the system perform, and to develop value-added functionality that differentiates one OEM's product from another. Unfortunately, with the complexity of network products today, the first two tasks consume so much effort that the task of adding value to the product is the last to get tackled. In the past, there has been, in a technical sense, a close coupling between all of these tasks that made the control-logic and system-performance tasks difficult to separate from the value-added functionality. New software technologies change all that.

Architecture for System Specialist Business Model

Distributed software technologies from the information technology (IT) world have wrestled with the problems of integrating independently developed applications running over diverse hardware and software platforms. Technologies such as interface definition languages (IDLs), code generation, object request brokers (ORBs), Java, extensible markup language (XML), and enterprise application integration (EAI) software frameworks emerged during the last few years to allow software applications to share data, share functionality, and bind stovepipe applications together into integrated business systems.

Companies such as Tesaria have begun to apply the same distributed software technology to the network equipment development world. At Tesaria, we realized that we would achieve greater developer productivity with this software bus architecture but that, more importantly, this software architecture supported the system specialist business model between Tesaria and our OEM partners.

In the system specialist business model, the system specialist produces fully integrated, network equipment compliant with the relevant standards. The OEM may further enhance the product with their own technology, but the product that ships from the system specialist may be used directly in operator networks.

Following the system specialist model, the Tesaria XGSN is a fully implemented GPRS support node (GSN). At the same time, we expose our internal, fine-grained, software-control functionality through our IDL so that our OEM partners can implement value-added functionality that supports their unique customer value proposition.

Let's look at how this works in practice. One of our partners plans to use content-transcoding capability to give end users faster perceived data throughput. The end users will be given individual control over the accelerator algorithms to suit their individual user tastes and preferences. This gives the OEM a key differentiator for their GPRS GSN product.

Tesaria supplies the GSN platform and enables the OEM to use internally developed, proprietary, content-transcoding technology where the OEM has a competitive advantage. By exchanging control messages between the Tesaria GSN core and the OEM's content-transcoding module, the OEM is able to deliver unique differentiation to its customers.

This is not a joint development project. Co-development on a single product-development project has a poor track record, as development teams and company management get in each other's way during the development. Tesaria has total control over the GSN core implementation, and the OEM partner has total control over the content-transcoding implementation. This is made possible by the software bus architecture and the IDL–defined programmatic interfaces between the two software systems.

The OEM controls the priority, timing, and effort involved in delivering their differentiating features. Concerns over sharing proprietary technology and trade secrets disappear because the OEM maintains full control throughout the development process.

How does this benefit the industry? All sectors of the GSM industry gain from this business model. GSM operators benefit from competition between OEMs, as the OEMs vie for operator contracts with product innovation and price incentives. Industry rationalization that results in OEMs exiting the business will leave higher prices and fewer product enhancements.

OEMs realize that if they exit the GSM equipment today, they would be foregoing a lucrative and strategic GSM equipment business that awaits them beyond the next year or so. Working with system specialist suppliers to manage their finances during this slowdown in network spending preserves their customer base and market position for a day when growth returns to this sector.

The system specialist business model sustains a competitive industry structure even as the OEMs share common overhead costs over a broader revenue base. Eliminating the wasted development of GSM standards-defined functionality allows the OEM suppliers to concentrate on creating value-added capabilities. The capabilities will give more services for GSM operators to sell to their customers and more reasons for the customers to subscribe to next-generation services.

Winning the Marathon

Competitors who know how to last the entire race win marathons. They know when to conserve energy, when to make a break for the finish line, and when to make strategy adjustments. The GSM equipment industry is clearly in transition. No one knows what the landscape will look like in three years' time, but it is safe to say that the industry will be more efficient than it is today. As the major equipment suppliers focus on their core business and find ways to cut R&D spending without losing competitive position, system specialist suppliers provide a means for the OEM suppliers to focus on product differentiation and to manage their bottom lines back to health.

Wireless Generations: Past and Future

Andrew J. Viterbi
President
The Viterbi Group LLC

The Standard Nomenclature and Its Misleading Implications

It has become commonplace to categorize advancements in wireless technologies by a generational label. Thus, the first analog phones of the early 1980s, constrained to a bandwidth of 30 kilohertz (kHz) or less, are designated as first-generation (1G). The actual breakthrough, though, was not in the means of transmission, which were simply the evolution of decades of frequency modulation (FM) technology, but in the cellular concept that, by assigning multiple base stations to the population of users, increases system capacity many-fold while reducing power requirements for the user terminals. Digital handsets, designed primarily for voice and standardized in the late 1980s and early 1990s, are called second-generation (2G). Lastly, advanced digital terminals, designed for both data and voice and standardized in the late 1990s, are designated third-generation (3G).

These simple distinctions may appear perfectly logical on the surface, but they do not reflect the marked difference between the two categories of 2G technologies: those based on relatively narrowband time division multiple access (TDMA) and those based on wideband spread spectrum or code division multiple access (CDMA). Even within the TDMA category, there are three varieties: on the one hand, digital advanced mobile phones system (D–AMPS) and personal digital communication (PDC), which preserve the bandwidth of analog technologies, and, on the other hand, Global System for Mobile Communications (GSM), which occupies a wider bandwidth and time-shares among a larger number of users. The relative success of the three TDMA varieties has been drastically different. While the more narrowband versions have reached their limits at only tens of millions of users and will gradually begin to be phased out, as has already largely happened with analog phones, GSM has proceeded to gain hundreds of millions of users, representing more than 80 percent of the total market. The latter's success was partly the result of early entry, though it owes more to European industrial policy, which froze out competitors but, by enabling universal roaming throughout the European Union, proved so desirable that it was extended to many other parts of the world. CDMA is a distant second among 2G technologies, with about 70 million users, mostly in North America, northern Asia, and parts of Latin America.

Yet the many demonstrated advantages of CDMA have become so widely recognized that it is the overwhelming choice for 3G, which is projected to handle equally both voice and data, the latter at speeds in the hundreds of kilobits per second (kbps) and, in some instances, as high as 2 megabits per second (Mbps). Even here, however, two rival approaches have been standardized—one championed by the current CDMA service providers and manufacturers, and the other by current TDMA entities. The former is called CDMA2000, while the latter is referred to as wideband CDMA (WCDMA), among several other designators for both. The differences between the two 3G CDMA standards are relatively minor, mostly small discrepancies in parameter choices, with one exception: the issue of whether or not to synchronize base stations or let them operate asynchronously. In the latter case, as proposed in the WCDMA standard, terminals must be resynchronized to the most recently encountered base station. The principal argument raised for asynchronous base-station operation is to avoid the current dependence on global positioning systems (GPSs) for synchronization. This seems strange for two reasons: first GPSs are not the only means to achieve this; other approaches to network synchronization have been successfully demonstrated. More importantly, a predominance of handsets, at least in North America, may soon incorporate GPS chips for emergency location purposes, making the issue of incorporating GPSs in the far smaller number of base sations of minor significance.

But the considerable difference between the two approaches to 3G lies not in the aforementioned details, but rather in the approach to the transition—designated as "2.5G technology." For the systems that currently employ 2G CDMA, the nomenclature is fitting as a way to describe the evolution from current low-speed data capabilities to the much higher rates anticipated by 3G applications. The required improvements are implemented mostly at baseband in signal processing chips, with virtually no changes in radio frequency (RF) front ends or antennas. They include more efficient modulation and demodulation of the current (direct sequence) spread-spectrum waveform, improved power-control techniques, and advanced error-correction codes suited to the less stringent delay restrictions of data. These are evolutionary modifications, which over time will lead to the full capabilities of 3G without major base-station refurbishment and with complete backward compatibility with 2G and 2.5G CDMA handsets. On the other hand, what is being referred to as the 2.5G transition for GSM (and possibly other TDMA formats) is really only an extension of GSM to handle higher data speeds. Known as general packet radio service (GPRS), it provides for multi-time-slot transmission of data, thus nominally reaching 64 kbps speeds, a minor milestone on the way to 3G requirements. More important, in no way does this represent a step toward implementing any form of CDMA. The goal of the GSM–standard group to implement WCDMA will require a complete refurbishment (or new construction) of base stations, at far greater cost and delayed time to market. Hence, the wide divergence in 2.5G approaches and the misleading implications of the term when applied to GSM.

Future Technologies

What newer technologies can we expect after the currently occurring transition to 3G is completed? The several candidates that present themselves all capitalize on one or both of the two primary characteristics of spread-spectrum technologies:

1. Universal wideband occupancy in each cell by many users, with advantages through frequency diversity, multipath resolution, and combining and interference management
2. Tight signal parameter control by exploiting the bidirectional nature of the communication channel and thus adapting both signal modulation and demodulation to the instantaneous channel conditions
3. Not mutually exclusive, multiple access technologies are potential contenders for what is already being called fourth generation (4G). These are as follows, in order of decreasing likelihood:
   - i. Multiple-input/multiple-output processing antennas, also described as spatial processing, space-time coding and "smart" antennas. It has for some time been recognized that spatial processing of multi-element antennas provides an added dimension for improvement beyond what is achieved by temporal processing. Some rudimentary forms of antenna-element processing are already present in some current generations. Rapid and accurate channel measurements achieved through spread-spectrum and bidirectional signaling, as mentioned, will benefit spatial as much as it does temporal processing.
   - ii. Orthogonal frequency division multiplexing and multiple access—a modified spread-spectrum approach that may provide for simplified processing and more rapid adaptation to channel conditions. In the forward direction (multiplexing), it substitutes orthogonal sinewaves for the Walsh functions of CDMA. In the reverse direction (multiple access), it has characteristics of both frequency division multiple access (FDMA) and spread spectrum.
   - iii. Ultra wideband (UWB) radios, which by transmitting subnanosecond pulses, somewhat randomized, effectively spread the spectrum over a bandwidth exceeding 1 Gigahertz (GHz). Because regulatory authorities do not award such wide swaths of bandwidth to any carrier or user, UWB proponents request that, due to the extremely low spectral density of such signals, they be allowed to coexist with present occupants, thus effectively rendering much of the presently allocated spectrum unlicensed—at least for UWB implementers. While the federal Communications Commission's (FCC's) action to initiate an inquiry on UWB technologies has represented a breakaway from existing mindsets, the outcome and wide adoption remains questionable. Also to be demonstrated are cost-effective terminal implementations. While time and major advances in enabling technologies may well overcome the physical uncertainties, the departure from established practice remains the greater hurdle.

Though measurable activity is occurring in each of the aforementioned, the types of compelling market needs and driving forces as had previously induced the overwhelming acceptance of 1G and 2G technologies have not yet appeared .

Section VI:

Operations and Quality Control

Centralized Network-Performance Monitoring

Gary Barton

Senior Market Manager, Mobile, Global Market Management

Agilent Technologies

Introduction

Providers can view performance monitoring in two ways: from the drive-test side and from the network side. This paper provides an overview of performance management from the network side.

First, it will explore why performance management is needed. Once the need for this is clearly established, issues of what performance management involves can be discussed. The two major monitoring techniques, real-time thresholding and historical reporting, will be introduced; and wireline, wireless, and Internet protocol (IP) performance metrics will be considered briefly. Finally, given that performance management is important for quality of service (QoS), and given the techniques, performance metrics, and architectures available for monitoring the network, this paper will look at how performance management can be implemented and integrated with the rest of the network-management tools.

The Need for Performance Management

Questions to Help Guide Providers

Service providers do not necessarily always know that they need performance management. What they more generally know is that they simply have certain problems that they need to address. Providers, then, face a number of questions, including the following:

- What is the perceived quality of the network?
- Is the network keeping pace with growth?
- Where should investment occur, and how much should be invested?
- How can more revenue be generated with the existing system?
- How can operating costs be reduced?
- How can a problem be identified as a simple abnormality or an indicator of something more serious?
- How can churn be reduced?
- How can future capital expenditures be predicted?
- How can system usage information be used to improve marketing and sales?

Along with these issues, recently introduced legislation has raised another concern. This legislation would require the Federal Communications Commission (FCC) to publish information on the number of complaints against a telecommunications provider and the nature of the problems. Under this legislation, the FCC would be charged with identifying responsible carriers. Thus, with the FCC reporting the number of complaints received about telecommunications networks, providers will also have to ask if they will be the next newspaper headline.

Definition of Performance Management

Performance management can provide the information necessary to answer all these questions. Performance management—defined as the collection, analysis, and presentation of network-performance data for use in evaluating the quality and efficiency of the telecommunications network—basically monitors the health of the overall network. It asks not only if the network is providing high-quality voice service but also if it is using all resources to their highest efficiency.

The service provider faces the challenge of striking a complex and delicate balance between quality at one end of the spectrum and growth at the other (see *Figure 1*). Providers will always look to add more subscribers, which, without a corresponding increase in the amount of network equipment, can result in increased utilization and degraded network quality. Adding more equipment will improve the quality of the network by reducing congestion, but overprovisioning a network increases the overall costs of a network. The goal for the provider, then, is to maintain the level of quality required to be competitive in the market without spending too much money in the infrastructure.

The Customer's Perception of the Network

Thus, a customer-centered point of view is critical to carrier success, and performance management provides that customer's perception of the network. Customer perception is especially important in two areas: purchase and retention. Among purchase considerations, price is the foremost in people's minds. Customers invariably will want to know how much they will have to pay for the phone and for service. The second-most important consumer-purchase consideration is coverage—the general geographical coverage area where good reception occurs. Services, such as voice mail,

FIGURE 1

Network Balancing Act

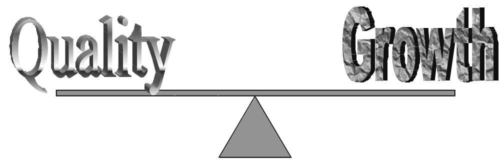

call waiting, and the like, make up the other area of customer concern. Customers are also highly concerned about availability—that is, the likelihood of finding a phone in stock at a store and walking out the door with it programmed and ready to use.

Customer-retention considerations add additional concerns to these purchasing considerations. Call quality, which includes voice quality as well as drops and blocks, are the most important for retention. Roaming service and customer service also have an impact upon retention decisions.

When data transmission is considered as well, more concerns emerge. The two major determinants of customer satisfaction in this case are the content of services (does a service provide value to the customer?) and the quality of the service (QoS). Quality of data services is based upon the provider's speed (which is essentially determined by the available bandwidth), jitter, and delay. The provider, for instance, must determine if there is sufficient bandwidth for Internet browsing, or if voice over IP (VoIP) or global positioning system (GPS) would experience too much delay and jitter within the system for customer satisfaction.

Figure 2 shows the results of a J.D. Powers survey concerning customer satisfaction. The survey found that after the consumer has purchased a phone and arranged for service, the most important factor for customer satisfaction becomes call quality. In this light, it should be noted that voice quality is only one piece of the overall call-quality perception; drops, blocks, and other call problems are the other parts that compose this perception.

Performance management, which can focus on these quality issues, then, has clear value in promoting favorable customer perception.

Components of Performance Management

Performance management is composed of two basic elements: real-time thresholding and historical reporting. A third area involves controls—the call gaps, filtering, and the like that manage the calls coming in once a problem is identified to manage overload in the network.

The first activity of a performance-management system is collection: each wireless network region involves gigabytes of data, making collection a nightmare. Storage is the next activity: once the data is collected, it must be stored in a database as quickly as possible. Next, the data must be managed: this process involves summarizing, as it would be impossible to query against just the raw data—there is simply too much of it. The data must then be analyzed to provide some meaning to the raw information collected. Finally, it must be presented in a format that is usable: reports (such as spreadsheets), graphs, or geographical displays. Many different kinds of users need to see the data collected in different ways, and the system must be able to supply what they need.

Real-Time Thresholding

Real-time thresholding is a technique for analyzing and presenting the data available from performance-management systems. *Figure 3* shows a scalable architecture for the real-time thresholding component. It involves multiple collection gateways that can serve the huge amount of data coming into the system. As the data is collected, the system applies a threshold to determine which data will be stored for analysis and which data will be discarded. In this way, data needed for specific purposes is maintained, whereas data that is not pertinent is discarded to avoid the need to store, summarize, and analyze it.

FIGURE 2

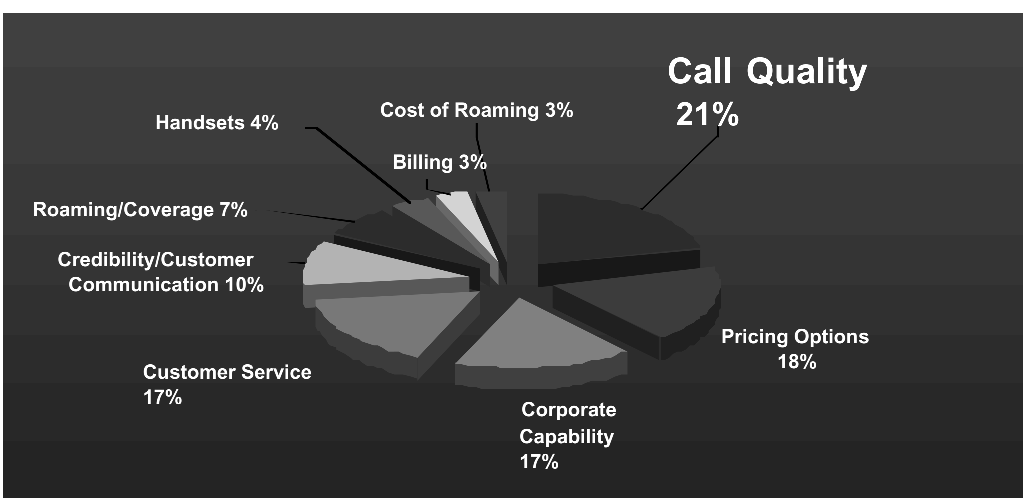

FIGURE 3

Real-Time Scalable Architecture

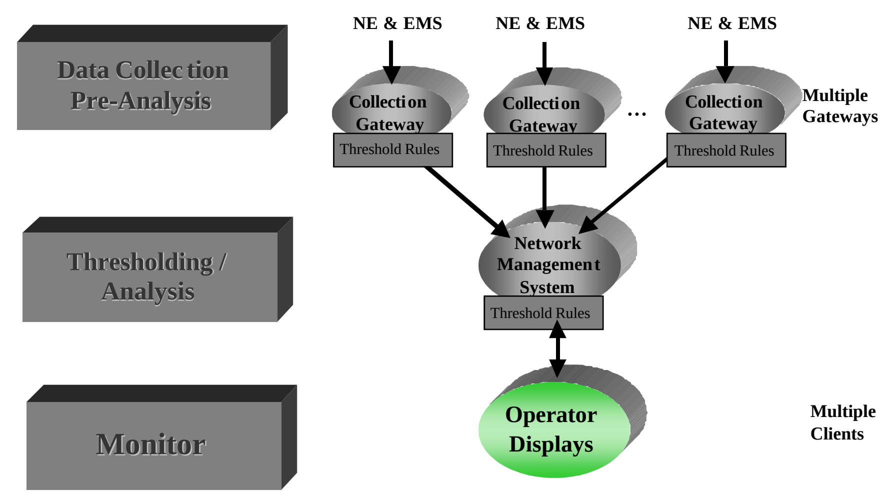

Performance Metrics

There are several ways to set the threshold or performance metrics of the performance-management system (see *Figure 4*). One method is to determine a static level, for example, to tell the system to issue an alert if a certain level is exceeded. Filters on the alert displays can then ensure that the proper person sees the alert. A static threshold is established by pre-establishing a specific value that the performance data is compared against. Two static thresholds can be established to form a range in which the data should stay within. Providers can set up groups of cell sites so that any new cell site in the group inherits the thresholding metric.

FIGURE 4

Real-Time Thresholding—Performance Threshold Metrics

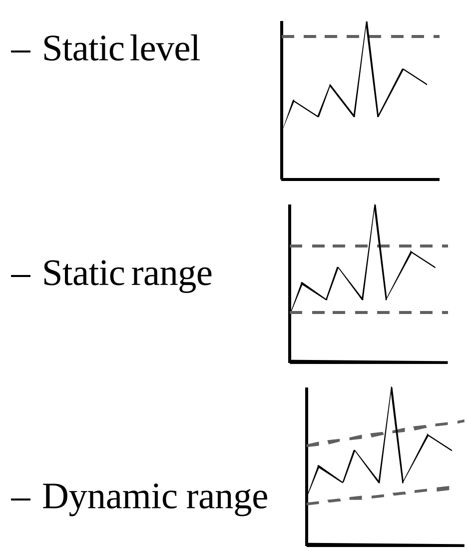

The advantage of using a dynamic range, on the other hand, is that the provider does not have to preassign the range that will generate alerts. New cells do not have to be told what their thresholds are; instead they learn them over time.

The performance metrics involve another important network task: normalizing the data. Networks today involve multiple vendors, multiple types of technology, and multiple versions of the same technology. To normalize data over this variety of components, generic metrics with vendor-specific, technology-specific, and version-specific formulas can be utilized. This approach allows the comparison of information from different types of networks (such as analog and code division multiple access [CDMA]) and different versions of the same software.

Historical Reporting

The historical-reporting aspect of performance management determines which data will be maintained for a longer time period. *Figure 5* shows a scalable architecture for this component.

As data is collected through the collection gateways (the average network will have multiple data loaders with multiple databases), it is stored in the database. The data manager reviews the data to determine if any data is missing. If so, data can be interpolated to fill the holes. From the raw data, summary tables are then developed (see *Figure 6*) to save space and decrease query time.

Network sizing is an important consideration at this juncture. The network must have sufficient bandwidth to pull data from far-reaching regions. Alternatively, the data can be collected and summarized at the regional level, then sent to a central collection point. Network sizing usually involves busy-hour considerations.

Historical reporting also is made up of several components (see *Figure 7*). The interactive report definition determines what reports will be issued, then saves the reports so they

FIGURE 5

Historical Scalable Architecture

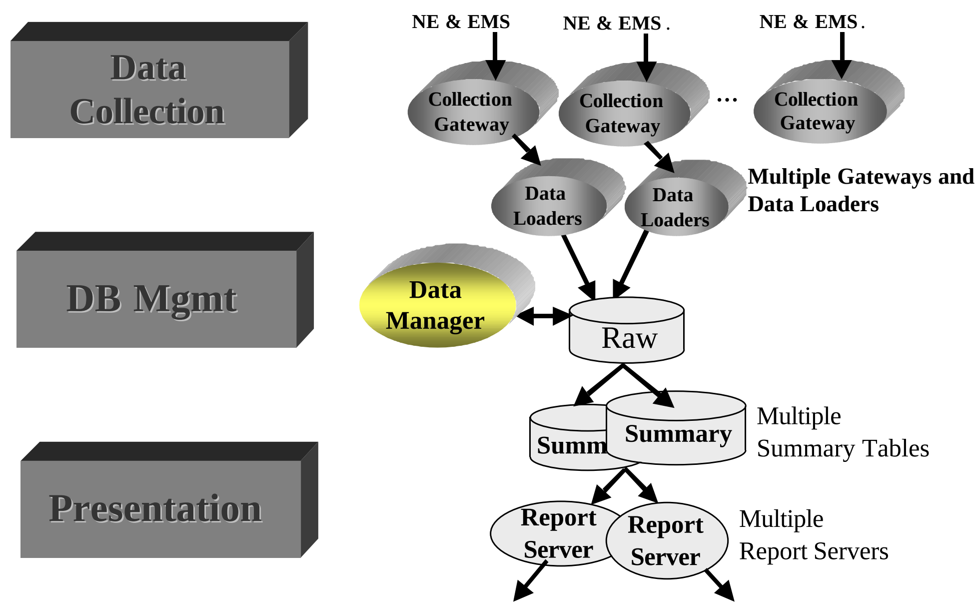

can be run again and again. Report scheduling specifies when reports will be run. For example, a provider may want to run a report every night to make it available to employees when they arrive at work in the morning. Historical reporting provides information about which reports are supposed to be run and when, and where the reports are stored after they are run. It can also handle interactive queries, ad hoc reports that are not required regularly. Engineers, for instance, may need to compare equipment functions over different time periods. Scheduled reports, on the other hand, automatically generate a spreadsheet or graphical display.

FIGURE 6

Summaries

- Used for saving space and query time
- Used in network sizing
  - Time-consistent busy hour
  - Bouncing busy hour
  - 3-day or 5-day average

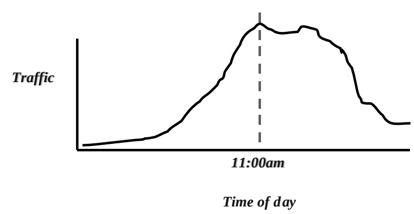

Performance-Management Implementation

Networks involve a wide range of users. At the top is the executive who makes decisions for the entire corporation. The executive typically looks for a simple number that can give some indication of how well the network is performing. On the other end are the engineers who are looking for answers to specific questions about the cause of specific problems. The performance-management system, then, must offer a reporting system that provides answers to the types of questions posed by the entire range of users.

Performance Reporting

The first requirement of such a system is a common reporting tool. The system should not have to build a whole new reporting infrastructure for each different user. The query tool should enable extensive reports, scheduling, and ad hoc queries across the company. Scheduling is another important aspect of the performance-management system, because it is much more efficient to set up the system to generate certain reports regularly than to request them each time. Web-browser access is critical for enabling anyone to access information from the system. Reports must be easy to modify so that information that engineers require to solve problems can be provided. Using a standard query language enables the company to hire people from the outside who can utilize the reporting system immediately.

From the historical reporting side, the system will ideally produce spreadsheets with a higher layer of summary report. This arrangement allows a user to see a problem, click on that area, and get additional information about it. This approach is much easier than having to run a series of reports (which requires an extended time period) to access the information needed. A graphical presentation is easier to use as well.

FIGURE 7

Historical Reporting

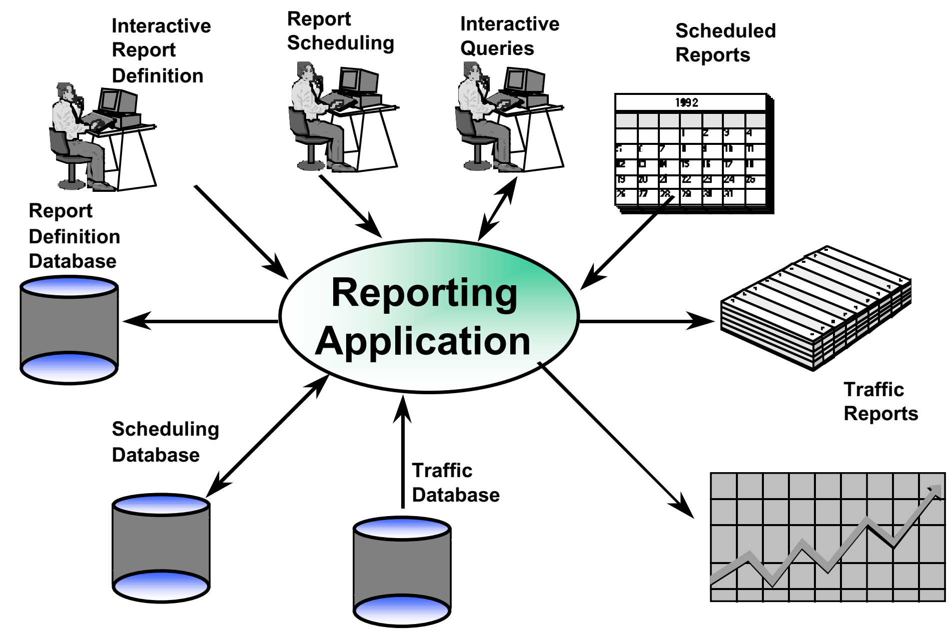

Historical reports are important in forecasting, especially for capacity planning. Exhaustion reports, which indicate when the system will run out of capacity, are one type of report used in this area. The other is a utilization report, which forecasts usage at a particular date in the future. This report can be used to plan purchases of equipment. A visualization display, including graphical or geographical (map-based) displays, indicates where failures or critical alerts are occurring and can also be used for fault management or trouble ticketing.

These performance reports are designed to monitor both grade of service and QoS. A grade of service is a predefined value that the network is engineered to: for example, a wireless network typically allows a specified level of blockage or overflow or a specified number of drops. QoS, on the other hand, is more subjective. It represents the customer's perception of service.

Wireless Metrics

Performance-management system metrics must monitor several areas. The most important of these is drops. This is a critical area because customers respond very negatively to being dropped from the network. Another version of a drop is a handover failure from one network to another. A handover failure appears to customers the same as a drop: that is, a customer was on the phone, then the connection was lost. Blocks, a situation in which the customer cannot access the network at all, would be the next most critical type of failure.

Capacity utilization must also be monitored to enable forecasting and equipment utilization. If the provider knows that equipment in one area is underutilized, it can be redeployed to another area where equipment is overutilized. This type of metric can also be used to identify cells that are "sleeping," or not carrying any traffic at all. These cells will not generate alarms, of course, because there is nothing to indicate that something is wrong with them. But when the provider finds no usage, a check may show that some problem has actually occurred, such as an antenna being blown over.

Wireline Metrics

Wireline metrics typically measure physical things, such as trunks and lines. These metrics can discern utilization of capacity. Overflows, or too much usage, may require rerouting traffic. Wireline metrics can also measure the total number of calls going through the network—information that can be very useful for forecasting and planning.

IP Metrics

IP metrics can present information about latency, which describes the delay in packets across the network; jitter, which involves packets reaching their destination out of order; and errors and lost packets. Lost packets require retransmissions, which consume additional bandwidth. And someone whose call was dropped will have to redial, thereby using up more network capacity. So savings can occur through performance management in these areas as well.

Centralizing Management

Performance-management metrics can fit into different types of management architectures: fully distributed, where each region is autonomous; fully centralized, where network responsibility is centrally located; or distributed with central monitoring, which has become most common today. Initially, local network management may involve a monitoring system for each of the different elements (see *Figure 8*). These element-monitoring systems feed information to a single network-management system. Alternatively, each different region may have its own local network-management system.

FIGURE 8

Local Network Management

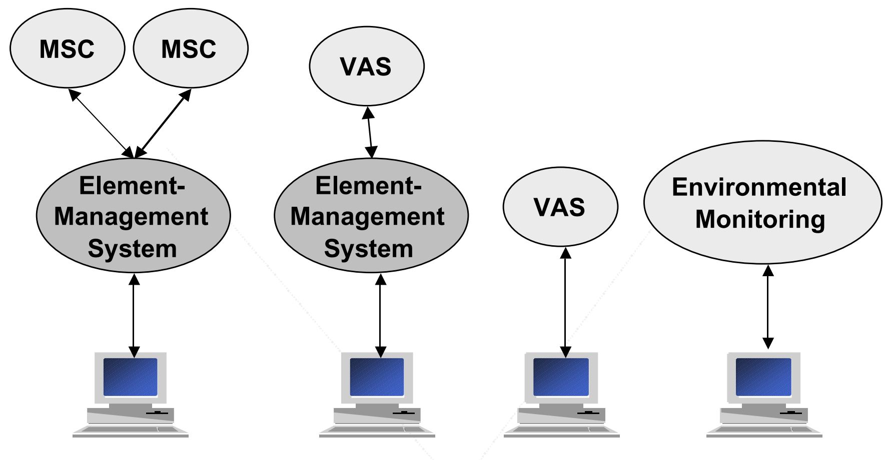

For security purposes, each region should be able to take over for the other regions in case of emergency. A central network-management system that can see all regions and take over management when necessary or when scheduled (for maintenance, for instance) may be the answer (see *Figure 9*).

Central network management, then (see *Figure 10*), looks at overall system performance across the entire network. In addition, it is responsible for identifying serious traffic problems, performing after-hours monitoring, and allocating resources networkwide.

Centralized management has several advantages. It supports a network-management strategy. Once the decision is made to utilize certain tools and maintain a certain level of quality, it implements that strategy across the entire network. Centralized management provides a common view across regions so that comparisons can be made. It allows consistent, standard alarm presentation to correlate activity across regions. A common platform integrated with common components expedites trouble ticketing. Common systems, interfaces, and practices reduce duplicate development.

In addition to offering advantages, however, centralized management also pose some challenges to providers. One such challenge comes when operating across multiple time zones (see *Figure 11*). Certain questions arise with this: What busy hour should be used? What time-stamp location should take precedence? Another challenge to operating across time zones is that equipment and operations are highly distributed: an operator can see equipment on a screen but cannot actually reach it physically. Low-bandwidth communication channels going across country can

FIGURE 9

Local Network Management—Centralized

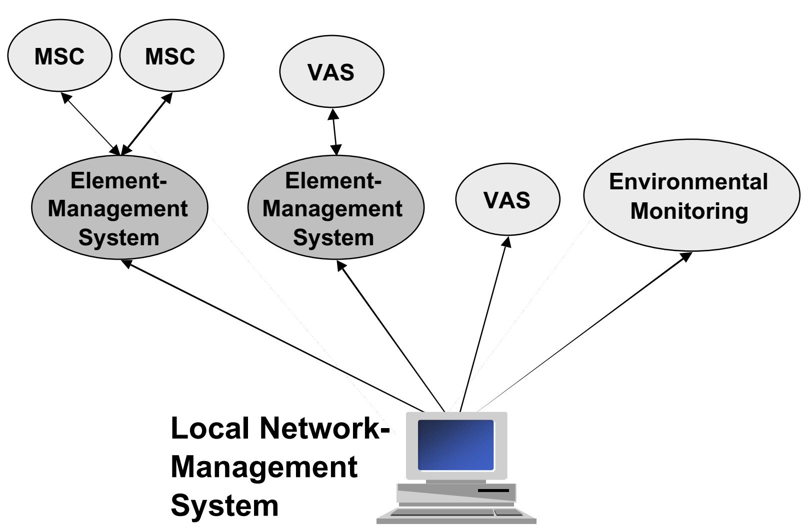

FIGURE 10

Central Network Management

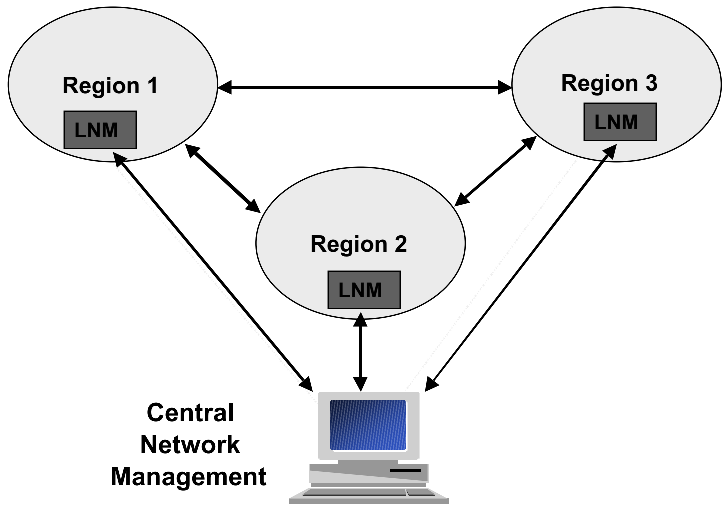

have an impact on company e-mail and performance data that must be transmitted from one coast to the other. When a national network is centralized, the network provider may find multiple versions of equipment and software as well. Getting everyone using the same version can be a daunting task. Finally, regions often want to maintain control themselves. They do not want "Big Brother" watching over them.

Integrating Performance Management

One of the most difficult aspects of any performance-management system may be integrating it into the company culture. Performance management, however, cannot be a stand-alone solution. It affects all the different areas of the organization. Performance management can function as a communication tool across the company, connecting such diverse areas as repair, facilities ordering, and marketing and sales. It can provide information about when repairs are necessary, when more equipment must be added, and where problems exist that cause customer complaints.

Conclusions

Performance management is critical to optimizing a network and reducing costs. It can ensure network health, enhance network utilization, support forecasting, aid in planning, and increase network availability through

FIGURE 11

Operating across Multiple Time Zones

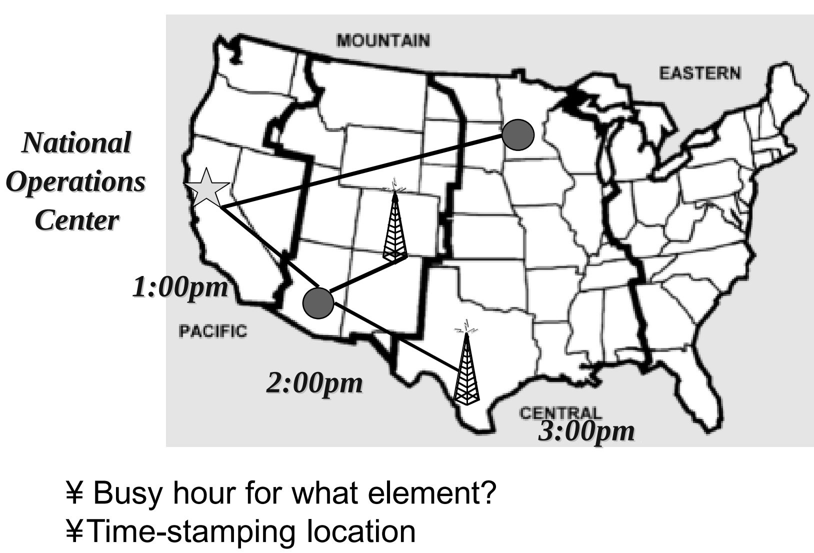

proactive management. Moreover, performance management can provide a view of the customer's experience with the network.

A good performance-management system must be implemented according to individual business needs, however. It must be determined, for instance, whether the business requires centralized monitoring, or if regions can be left to work on their own. Internally, the system must meet specific needs as well. It must consider all the different users and the kind of reports they need to do their jobs effectively. Finally, the performance-management system must be fluidly integrated within the organization: the data such a system provides is very valuable, and it should be available to and used by everyone in the company.

The OSS Architecture Crossroads

Silo-Based OSS Architecture Has Led the Industry up a Blind Alley—a Truly Integrated Approach Is Now Required

Peter Briscoe

Principal Consultant

Cramer Systems

A system that automatically provisions new telecom services is a holy grail for all operators. Faced with increasingly complex networks and a worldwide shortage of skilled provisioning engineers, ever more powerful provisioning software is the only hope that telcos have of cutting service-delivery times and improving network utilization.

Yet this goal may never be realized unless operations support system (OSS) vendors and telcos abandon the piecemeal approach that has characterized the development of their OSS systems. This has led to the so-called stovepipe or silo-architectured OSS suite in near universal use today.

Silo-based architectures have been the quick fixes for each new technology that has come along. The downside for telcos is that their network information has, as a result, been scattered across several applications, saddling them with a fragmented picture of their network resources.

The way forward is to adopt a more centralized OSS architecture, which holds network information in a central repository and makes it available to all OSS applications that can be used to obtain a detailed overview of network infrastructure.

What Is Silo Architecture?

In a stovepipe or silo-architectured OSS suite, provisioning for different technologies is supported by dedicated, standalone applications, and each is optimized for that function alone.

For example, a provisioning engineer might use one application to allocate capacity on a synchronous optical network (SONET)/synchronous digital hierarchy (SDH) high-speed ring and a second application to configure an Internet protocol (IP) or asynchronous transfer mode (ATM) layer in order to set up a new circuit for a customer. Separately, each has access to the telco's order capture software and to the network activation layer.

The OSS architecture typically comprises a series of provisioning stacks connecting the customer order-capture layer with the network—a structure that is suggestive of a grain silo or factory stacks, hence the name. These provisioning applications are generally standalone, but communications between applications can be a serious problem.

How did this situation arise? Until fairly recently, switched voice services accounted for most of the telecom market. The provisioning systems that supported them were developed internally and were tailored to suit each company's process and needs.

However, explosive growth in data traffic in the 1980s, underpinned by a succession of new technologies such as SONET, ATM, transmission control protocol (TCP)/IP, and integrated services digital network (ISDN), was beyond the scope of most in-house OSS teams to support. Instead, telcos turned to external vendors of provisioning software and bought dedicated OSS applications supporting a specific technology.

These systems were bolted onto the telco's existing OSS applications suite without any major reworking of the system, thus minimizing software-development costs. Although they were designed for the job and worked well, there was little integration between the new application modules and the core OSS software. In particular, there was no inherent mechanism to exchange data. This proved to be a major problem.

Changes to the OSS software were often matched by organizational changes. Many telcos set up a new network provisioning team to support each new service. This, in turn, encouraged a fragmented view of the network and stored up trouble for the future.

Consolidation across the telecom sector has, ironically, further exacerbated the OSS problems for some telcos. Takeover of one telco by another often meant that an OSS support team was left with a legacy of two incompatible OSS systems. More often than not, the merged company chose to keep both systems running in parallel rather than migrate network information to one system. This decision usually prevented the telco from realizing the full benefits of consolidation and network rationalization.

The Communications Bus

One fundamental problem with a silo-based architecture is that it presents a fragmented view of the network and fails to take account of the inter-relationship between individual provisioning operations, such as network routing and checking quality of service (QoS) levels when planning a new end-to-end connection. Each activity is generally handled by separate applications. For example, there are usually dedicated silo stacks for provisioning asymmetric DSL (ADSL) connections, IP routers, SONET optical fiber buses, and dense wavelength division multiplexing (DWDM) connections.

In-house OSS support departments have attempted to solve these communications problems by building links between applications. But such was the complexity associated with many point-to-point connections that it soon became clear that a formal method for exchanging distributed information between modules would be needed. The enterprise application interface (EAI), such as TIBCO Software, has offered one attractive solution.

Several companies tried to apply this new EAI technology with the next stovepipe solution to be added to their OSS suite, but it soon became obvious that all the system components would need to be modified. Each would need an EAI communications adapter for the exchange information via a common EAI information bus.

As a result, this proved a far bigger problem to solve than originally envisaged. With no prospect of an immediate financial payback, the idea of adding an EAI information bus across the entire OSS suite was often shelved. Many telcos have since recognized the crucial strategic role played by OSS in network profitability and customer service and are revisiting the case for a major re-engineering of their OSSs.

Toward a Centralized Data Model

In a silo-based OSS architecture, network information is fragmented between applications. To make this information available to other applications, the industry is moving toward the adoption of a complex messaging technology. But a better solution would be to centralize network information in a common database model of the network.

Some equipment vendors and system integrators have adopted technologies such as lightweight directory access protocol (LDAP) and the common information model (CIM) (sponsored by the Object Management Group) to provide a common framework for the sharing of network information. But these central store network models used a broad-based generic description of network components to ensure that all the systems involved with the OSS could use it.

Databases built on this model are complex and difficult to use, so systems integrators developing centralized OSSs have, as a first step, arranged for the network model to be updated periodically by uploading data from the OSS applications on a batch basis.

These centralized data models provide greater visibility of the network and can be used by provisioning engineers to guide their routing decisions. But the network information provided is read-only and does not allow provisioning engineers to interact with the network model in order to allocate network resources. Consequently, it fails to address the main problem with silo-based systems—the lack of end-to-end control over the provisioning process.

This is because the task of allocating network capacity is divided between several provisioning engineers, each of whom is responsible for provisioning a specific network layer or technology (see *Figure 1*). Individual provisioning engineers can only see their network layer and cannot examine the impact of their provisioning requests on other network layers. These requests are passed on to their colleagues to handle. With the addition of each new system stack, the production line just gets longer and the provisioning problem gets bigger.

Case for Network Overview

For a provisioning engineer to do his job well, he should have access to an end-to-end view of the network showing where capacity has been allocated and where there is capacity to spare. To illustrate the problems that can arise, a good example would be the task of adding a new virtual private network (VPN) customer to an existing network.

The first step is to locate the nearest access points to the customer's sites. Once you have checked that capacity is available, a path is routed across the access network to a point-of-presence (POP). This will provide access to the telco's backbone network that may typically be a combination of ATM and optical technologies.

Selection of the nearest access point is a simple matter. The provisioning engineer has only to check the status of the nearest access point to see what ports and slots are free. But this check says nothing about the connection between the access point and the POP or across the backbone. This information is normally held in the IP, ATM, and optical applications.

An optimal route may be found to the major POP, but there is no way of knowing that the connection into the network over the next layer has capacity. *Figure 1* illustrates the problem.

A customer request for a new service is mapped on to the IP layer first and then onto the adjoining ATM layer. Separate engineering teams for the IP and ATM layers might typically handle this process. *Figure 1* shows that the IP POP maps directly to the underlying ATM switches. The selected POP has capacity, but the ATM links from the switch are full. So the ATM planner either has to reject the request for a link and hand the routing task back to the IP planning group or must arrange to add capacity on the requested path.

A cross-technology view of the network avoids this problem by presenting the provisioning engineer with a view of network capacity across technology layers. This means he can explore alternative routes and get the routing problem right the first time. For example, when a route through the first POP and ATM backbone could not be found, the engineer could have selected a second POP and used the available ATM backbone links.

By contrast, in a silo system, the routing operation has to be repeated until a free route though the ATM layer can be found. This presents a fundamental barrier to the further automation of the provisioning process. After all, what is the point of automating a process that may have to be repeated several times?

FIGURE 1

A Silo-Based OSS Architecture

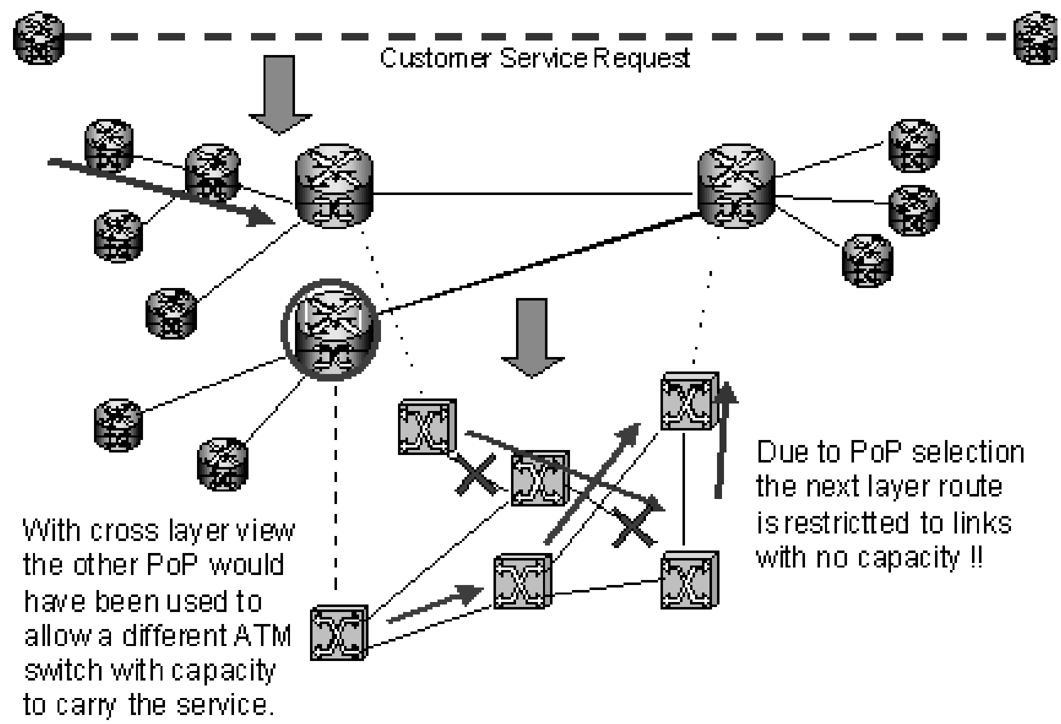

Central View Importance

An OSS, which presents a single view of the network at a physical and logical level, is fundamental to any hope of automating the network-provisioning task.

But implementing a central view can still be a challenge, because its effectiveness will depend on the accuracy of the data making up the network model. This information can be dispersed throughout an organization. It may be held in different data formats; some of it may even be paper-based. The task of collating data and loading key parameters can be crucial to the success of any project of this sort.

Unlike a central data view, the structure of the data is tailored to the functions that must be supported by the OSS, but access to this data must be provided in an open and easy fashion for read-only and reference purposes, otherwise it will revert back to a closed system approach.

The goal for many service providers is to balance the transfer of functionality from the existing systems to the new central view, as the new view provides a significant business benefit. For example, move the planning of ATM to the central view when the view also allows the mapping of passive devices and SDH/SONET usage.

The move to a central system is not simple, but it is a necessary step to allow today's operators to enhance their ability to speed the delivery of new services and ensure the fastest time to market—the keys to good market growth.

This article was originally published in "Telecommunications International" in July 2001.

OSS Integration Approaches

Lee Chow
Senior Analyst, Telecom OS
RHK, Inc.

Introduction

This paper addresses the various operations support system (OSS) integration and deployment approaches that are being taken in the telecommunications industry today. It will examine OSS interoperability needs and deployment approaches, including management suites, interface integration, and systems integration of commercial off-the-shelf (COTS) components. Also under consideration are industry trends that may affect each of these approaches and may preclude or enable the future proliferation of some of these approaches.

OSS Interoperability Needs

Figure 1 describes the framework of a typical service provider's business processes. One or more of these processes are instantiated and supported in OSSs, including billing, customer-care, service-provisioning, and service-activation platforms. These OSSs support the service providers' business processes and, when deployed, are tightly integrated. The fundamental reason to integrate the processes is to ensure process flow-through and, as a consequence, enable automation. Service providers are looking for automation to derive profitability and also to eliminate manual processes that are cost-prohibitive. An example of process flow-through is service fulfillment, where the order of handling, service-provisioning, network-provisioning, and element-provisioning processes are tightly integrated.

Management Suites

OSSs today are deployed essentially using one of three approaches. The first approach is via management suites. Two examples of the suites being offered today are shown in *Figure 2*. Offered by a single vendor, management suites are pre-integrated, shrink-wrapped solutions. These solutions span horizontally across the fulfillment, assurance, and billing space and also vertically across the telecommunications management network (TMN) layers. These management suites will have applications in the element-management layer (EMS), network-management layer (NML), service-management layer (SML), business-management layer (BML), and also in the fulfillment, assurance, and billing space.

The components of the management suites are tightly integrated, most likely through proprietary interfaces. For example, with Lucent One Vision, there is a billing component, a service-activation component, and a trouble-ticketing component, all tightly integrated through a proprietary Lucent interface. Because of the nature of the proprietary interface, it does not lend itself well to a multivendor OSS. For example, if the billing component of a management suite is offered by another vendor, integration between the multivendor components will present a challenge, primarily due to the proprietary interface. Mulitvendor integration will more likely occur at the OSS–to–network element (NE) layer and also the OSS–to–element-management system (EMS) layer.

OSS Integration through Management Interfaces

The second deployment approach is through OSS interfaces (see *Figure 3*). OSS interfaces can be applied to several reference points. The requirements, approaches, and solutions vary depending on the reference point. For the interfaces between EMSs and NEs, there are many NE view models and protocol-specific implementations that exist today, including simple network management protocol (SNMP) and transaction language–1 (TL1).

Standard information models and interfaces exist today but apply only to technology-specific domains. The standards include the recommendations from the Asynchronous Transfer Mode Forum (ATMF) and the Synchronous Optical Network Interoperability Forum (SIF). The majority of the interfaces between the EMS and NE are proprietary, driven by the fact that many of the equipment vendors who provide NEs also provide EMSs. When the vendor's bundled NE and EMS are deployed in a service-provider environment, that NE–EMS interface is transparent to the service provider. Therefore, the proprietary nature of the NE–EMS interface does not present an issue for service providers. Also, EMSs typically have a single-vendor focus and do not normally apply to multivendor network elements

At the reference point between the EMS and network-management system (NMS), there are many network view models that exist as well as many different types of interfaces, such as TL1, SNMP, common object request broker architecture (CORBA), extensible markup language (XML). Some of these solutions are recommended by the International Telecommunication Union (ITU). Also noted are T1M1 submissions to ITU, the ATMFM4 interface, and

FIGURE 1

OSS Interoperability Needs

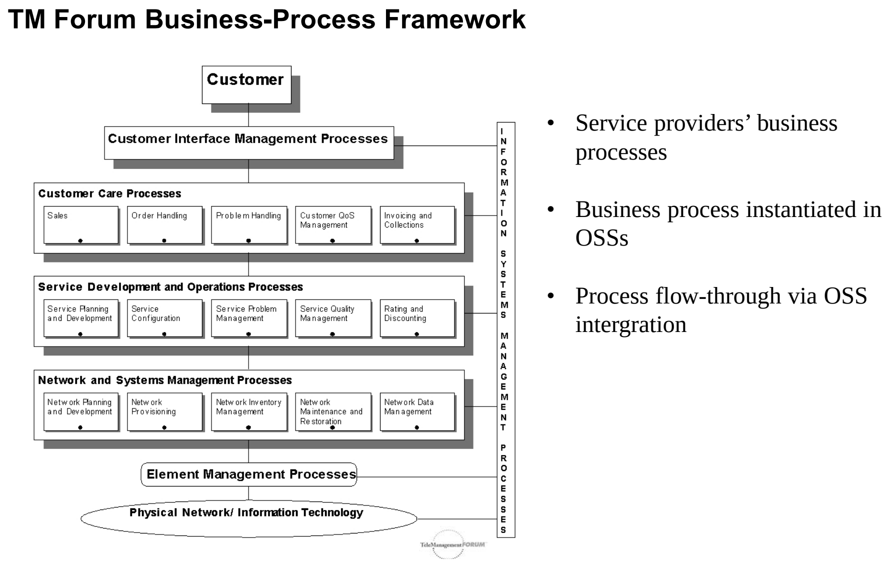

the SIF interface. Work within the TeleManagement Forum (TMF) has produced the TMF 509 interface, which is a CORBA–based interface applicable to SONET, synchronous digital hierarchy (SDH), dense wavelength division multiplexing (DWDM), and ATM.

Proprietary interfaces are used between the service-management and network-management layers. Work in this area, also in the TMF, is underway under the TMF508 initiatives.

The approaches in multivendor OSS integration are accomplished today through proprietary interfaces, standard recommendations, or through de facto standard interfaces. De facto standard interfaces are interfaces that are not formally or fully recognized and yet have enough momentum from the vendor community that they proliferate through the industry.

The interfaces used for OSS integration are segmented into several different components. Functional requirements and use cases are delineated from the information and data model and are further differentiated from the technology-specific implementation interface. The motivation for the segmentation is to drive the preservation of

FIGURE 2

Management Suites

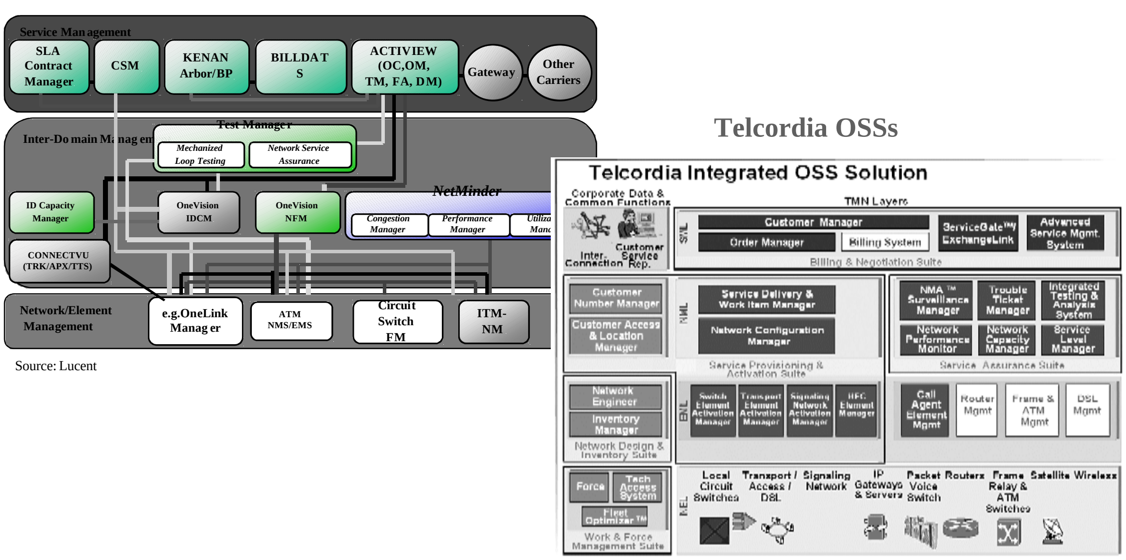

Source: Lucent

Source: Telcordia

FIGURE 3

OSS Interface Integration

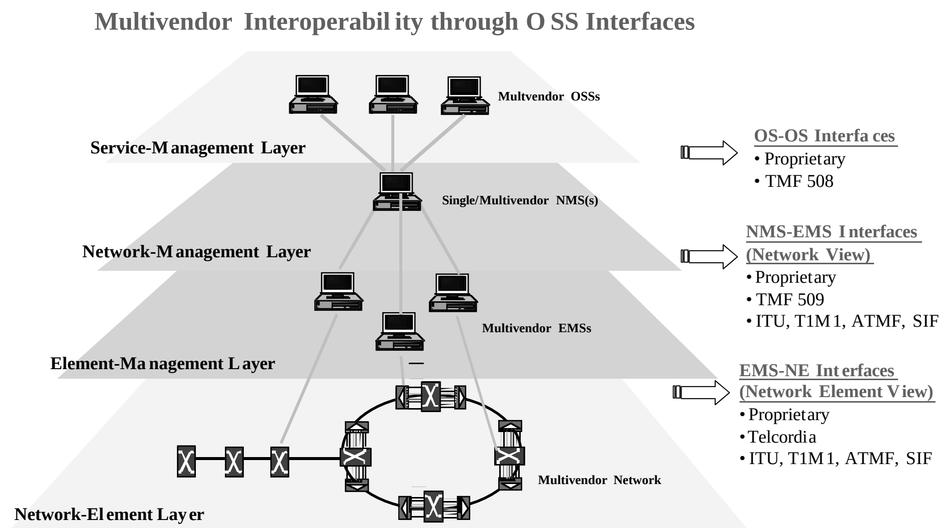

the data modeling effort, separate from the technology-specific implementation.

Within the interface-integration approach, vendor product compliance does not ensure multivendor interoperability. Multivendor interoperability is insured only after there is some level of interoperability test. Although vendors' products may comply to interface recommendations, there is still some room for interpretation. The onus of validating interoperability is, thus, left to the service providers, who will actually perform the interoperability tests.

Systems Integration of COTS Components

Systems integration of COTS components represents the most recent approach in OSS integration. The fundamental principle behind this third approach is a distributed architecture composed of interchangeable plug and play (PnP) components and integrated via a common communications bus. The multivendor PnP components interface and communicate through a common communication bus and reference data through a common information model. A workflow engine is used to coordinate the flow of the service provider's business processes.

An important aspect of this third approach is that interoperability is tested and ensured by third-party systems integrators (SI). The total solution is pre-integrated by the SIs and delivered to the service provider. Unlike the previous OSS integration approach, the integration interoperability testing is done not by a service provider but by the SI.

Figure 4 illustrates this particular approach, using the TMF Broadband-Cable Catalyst Project as an example. The PnP components include an order-management component from Metasolv; a billing component from Portal; a service-level agreement (SLA) management piece from Trendium; a performance management (PM) part from TSCI; a fault management (FM) component from Micromuse; and a service-activation piece from Syndesis. Each of the multivendor PnP components communicate through a common communication bus. If a common information model is not implemented, a data-transformation application is used to normalize the data representation. Finally, a workflow engine is used to enforce the service provider's business flows.

OSS Deployment Approaches

In examining the three approaches, two common factors of each approach become critical for service providers: time-to-deploy and the level of multivendor interoperability (see *Figure 5*). The management suites, because they are pre-integrated and are shrink-wrapped, will have a shorter deployment span. On the other side of the spectrum is integration through OSS interfaces, where there is a comparably higher time to deployment. In between is the systems-integration approach. In terms of multivendor support, management suites are very single-vendor focused. On the other hand, the systems-integration approach fosters a multivendor environment.

Industry Trends

The OSS integration and deployment approaches are greatly affected by emerging industry trends. One trend that is having an impact upon the systems-integration approach is being fueled by the greenfield carriers. Greenfield carriers are emerging carriers, typically without legacy OSSs in place with stringent time-to-market requirements. OSS deployment cycles are of primary importance to these service providers. Many do not have the resources to support in-house development, interoperability testing, and integration. As a result, these service providers will lean on an SI to do some of that work. The systems-integration approach is much more appealing to greenfield carriers.

FIGURE 4

Systems Integration of COTS Components

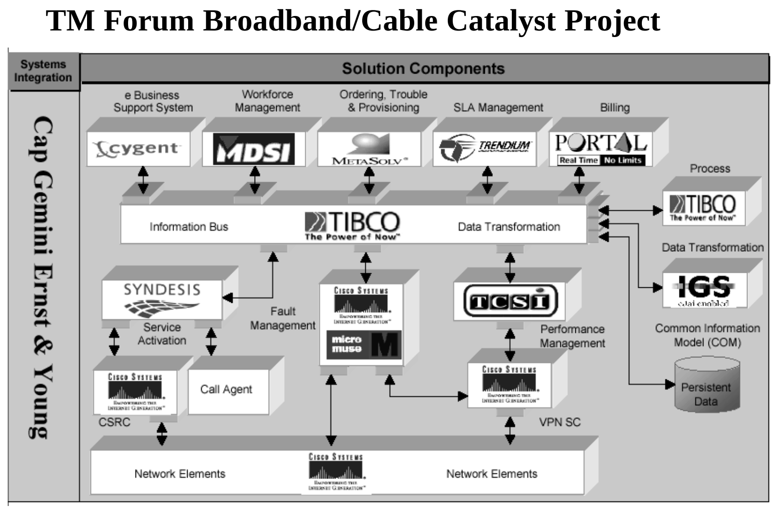

Source: TM Forum

Another important trend to note is the tremendous consolidation that has occurred in the OSS industry. Examples of such consolidation include ADC's acquisition of CommTech, Orchestream's takeover of CrossKeys, Agilent's purchase of OSI, Spirent's acquisition of Hekimian, and Nortel's acquisition of Architel. The trends toward consolidation may lead to the monolithic management suite approach. Prior to an acquisition, an OSS vendor's product may address a particular area, such as PM; after a complimentary acquisition, however, that product offering will likely expand from PM to FM to SLA management.

Conclusion

This paper examined three different OSS deployment approaches: management suites, systems integration of COTS components, and OSS interface integration. The management suites that exist today are single-vendor focused, shrink-wrapped and pre-integrated. OSS integration through interfaces utilizes three types of interfaces: proprietary, de facto, and standard recommendations with proprietary extensions. Finally, the last approach follows the fundamental principles of a distributed architecture, PnP components and integration through a common communications bus and data model.

FIGURE 5

OSS Deployment Approaches

Comparison of Approaches

(Time versus Interoperability)

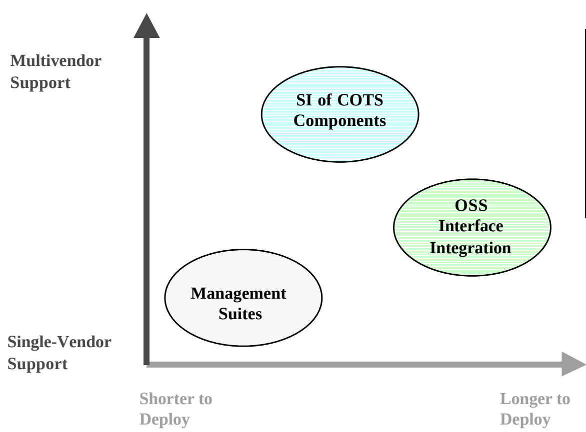

	Time-to-Deploy	Multivendor Interoperability
Management Suites	Low	Low
SI of COTS Components	Medium	High
OSS Interface Integration	High	Medium

QoS Routing Hierarchy in ATM Networks

I-Ping Chu
Associate Professor, CTI
DePaul University

Yuxing Tian
Senior Software Engineer, iDEN Division, GTSS
Motorola

Abstract

Quality of service (QoS) routing enables network service providers (NSPs) to guarantee a certain level of network performance to user applications while optimizing the network resources utilization. In this paper, we present the state-of-the-art of the process of constructing asynchronous transfer mode (ATM) private network-to-network interface (PNNI) hierarchy that enables the QoS routing. A new concept of inner metric to the PNNI logical group node has been introduced that enables clustering algorithms that recursively partition a network into peer groups and that improves algorithms to transform a generic peer group into a Complex Node Representation with minimal error. The ATM PNNI hierarchy created with these improvements will enable the network connection admission control (CAC) to build QoS routes and guarantee end-to-end performance requirements of user applications.

Introduction

In recent years, the Internet has created a huge impact on people's lives. It demonstrated the power of information sharing and the potentials that may bring. However, as the nature of the Internet protocol, namely transmission control protocol (TCP)/Internet protocol (IP), is best effort, when the information carried in packets are routed through different paths to reach a destination, there is no performance guarantee from NSPs. As a result, a broad range of network applications is failed by the current TCP/IP technology. This class of applications ranges from real-time audio/video and teleconferencing to bandwidth on demand, etc., and can be identified as applications with QoS requirements.

QoS is in general measured by a set of parameters characterizing connection performance. The most important QoS parameters are usually in terms of network resources and connection quality, such as central processing unit (CPU) time, delay (transmission and queueing), link bandwidth, and loss ratio, etc. A QoS guarantee is a performance contract between a user application and network, which is often in terms of end-to-end measurement. Asynchronous transfer mode (ATM) technology has been deployed to support this class of QoS–sensitive communication applications.

ATM is a connection-oriented technology. The use of fixed-size packet (called cell) enables the switching technology to route data packets at extremely high speed. It has already been deployed by the current Internet to provide *flat* Internet backbone networks. With its routing hierarchy built to the ATM PNNI standard, it provides an efficient routing protocol with excellent scalability and security for large networks

QoS routing along with other network-management functions, such as CAC, enables NSPs to provide QoS guarantees to QoS–sensitive applications. It uses connection-oriented techniques to make resource reservations at the connection-admission phase and guarantees its performance commitment. The data transmission of such connections will not affect or be affected by traffic dynamics of other connections sharing the common links. Before the reservation is done, the routing function must select a path with the best chance to satisfy the resource requirement of the new QoS connection request. Admission control determines whether a connection request should be accepted or rejected based on the required QoS and routing hierarchy it maintains. Once the request is accepted, the required resources must be guaranteed.

ATM PNNI Routing

PNNI is a routing hierarchy that partitions a network into a hierarchy of sub-networks, called peer groups, and aggregates network-topology state information. As a source routing protocol, each switching node needs to maintain a complete up-to-date network-topology and link-state information. Because of the size of the Internet, it becomes necessary for the network-state information to be summarized and distributed efficiently to minimize the memory and speed requirements on network switching nodes and to reduce network-management complexity and its overhead consumption of the network capacity.

A peer group can hide the details of its internal structure from the outside world by advertising a summarized view of the internal structure, hence providing excellent scalability and security measures. In the PNNI standard, each peer group is represented by a logical group node as an abstraction in the next-higher routing hierarchical level. The logical group node is represented by the PNNI Complex Node

Rrpresentation that retains the most important link-state and topology information of the peer group. In each peer group, a leader aggregates and distributes information for maintaining the PNNI hierarchy. The PNNI routing hierarchy also allows asymmetry in the sense that, for a given lower-level peer group, its parent peer group can be a grandparent or great-grandparent peer group to some other same lower-level peer group at the same time. An example of a complete PNNI hierarchically configured network is shown in *Figure 1*.

Routing collects and maintains network-state information and searches state information for a feasible path. The search for feasible QoS paths greatly depends on how the state information is collected and where the information is stored. The ATM PNNI routing hierarchy provides a standard structure to collect, aggregate, and store such state information. It can be viewed as a recursive process to generate peer groups. In the process of this abstraction, information about a child peer group of a logical group node is aggregated and summarized into a logical group node, which in turn is represented by the Complex Node Representation.

The PNNI is a QoS–sensitive source routing protocol. When a new connection request comes to a network-entering switch, the switch is responsible to find a connection path to reach its destination while satisfying its QoS requirements. The CAC function used by this originating switch will select a route with the link-state information of the whole network in its local topology database. If a path can satisfy the QoS requirements with the information maintained in the local link-state database, the path is encoded in a designated transit list (DTL). Then a message is sent out to set up the new connection and make necessary network resources reservations along the path. If, at a certain point, the connection's QoS requirements cannot be satisfied, a crankback occurs to partially release a connection set-up and comes back to the node that created the DTL to perform alternate routing. This probing process continues until either a path is found satisfying the QoS requirements or the connection request is rejected. After a successful attempt, a new virtual connection is set up; the source may start communicating with the destination; and the network guarantees its resource commitment to the QoS requirements for the lifetime of the connection. In a network, there may be several paths between two border nodes. The path chosen is always the one with the minimum end-to-end delay.

The PNNI uses flooding as its advertising mechanism. It ensures that every node in a peer group has an identical topology database. Furthermore, the flooding flows downward into and through the child peer groups. The PNNI routing hierarchy can be naturally viewed as a recursive generation of peer groups. At the lowest level of this PNNI hierarchy, clusters of the physical switches are grouped into peer groups. Then the detailed nodal and link-state information of each peer group is aggregated into one logical group node to a higher level in the hierarchy. This grouping and aggregation process is repeatedly applied to the higher levels of the logical group nodes in the PNNI hierarchy until the whole network is aggregated into one peer group.

PNNI Complex Node Representation

To be able to make effective QoS routing decisions, the most important topology-state information of a peer group needs to be retained in its logical group node. The complex node representation was proposed to represent the aggregated state parameters of the peer group in a logical group node. A node identifier identifies a logical group node. The complex node representation aggregates a peer-group topology into a star topology whose center serves as the interior reference point of the logical node and is called the virtual nucleus of the complex node [1]. The links connect the nucleus only to the border nodes of the peer group. This Complex Node Representation reduces the overall complexity and hides the topology internals of the peer groups. The logical connectivity between the nucleus and a border node is referred to as a *spoke*. The concatenation of two spokes represents a traversal of the peer group. This process of

FIGURE 1

A PNNI Hierarchical Network

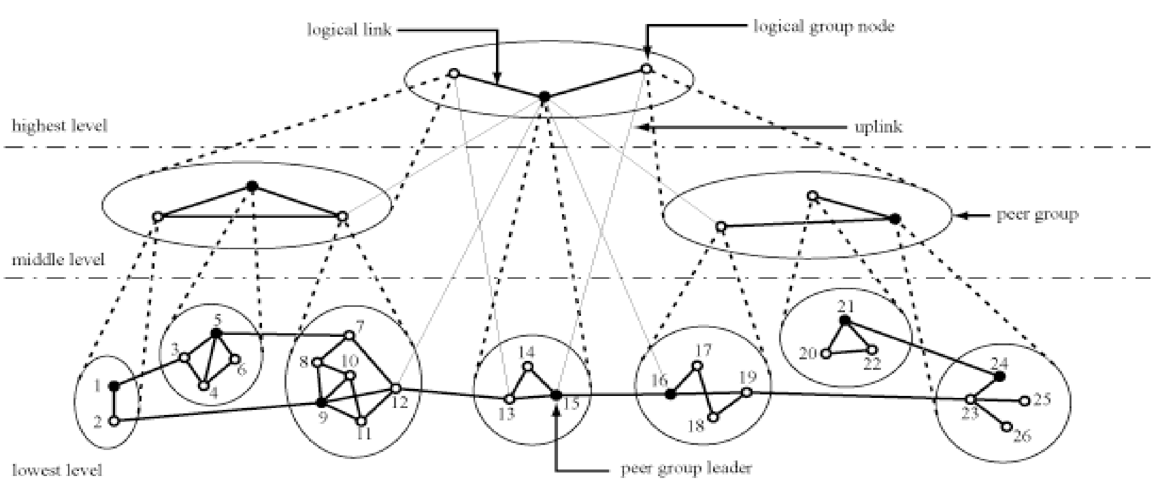

aggregation and summarization is called topology aggregation. As peer groups are usually not symmetric, exceptions are used to represent particular ports whose connectivity to the nucleus is significantly different from that of the default, or between two ports, i.e., a bypass that is significantly better than traversing the nucleus.

Many efforts for constructing effective PNNI Complex Node Representation with respect to various QoS requirements are discussed in various literature [3, 4, 7, 8]. In a typical aggregation model, a PNNI peer group is represented by a graph G. Let G = (V, E, B), where V is the set of switching nodes, $B \subseteq V$ is the set of border nodes, and E is a set of bidirectional links among the nodes in V whose values represent the corresponding link costs. This model can be applied to the networks where the majority of data traffic is balanced, as those in voice or video telephone conversations, etc., where the traffic volumes in both directions are about the same. A border node is a logical node that has at least one link across the peer-group boundary that it is in.

An N x N matrix is used to represent the cost metric between border nodes of a peer group, where $N = |B|$, which is the size of the border node set. Many aggregation techniques in the literature [2, 3, 4, 8, 10, 16] extensively use matrices to hold the topology metrics of concern. In [16], a new concept of intra-nodal metric is introduced that uses the main diagonal of the metric matrix $a_{i,i}$ to hold the intra-nodal costs, i.e., the default spoke value of the PNNI Complex Node Representation for the logical group node. Note that, at the lowest level of the PNNI hierarchy, logical nodes are real physical switches in the network. We initialize this intra-nodal cost metric to zero for the lowest-level nodes for the negligible network costs, such as delay, inside a physical switch. But at all other levels, this intra-nodal metric is obtained by the default spoke of the logical group node, and it is always a non-negative number. This internal metric enables us to construct a PNNI complex node for any peer groups with minimal error, which significantly improves the quality of peer-group aggregations and minimizes the discrepancy between the PNNI Complex Node Representation and the peer group it represents. The links in E are physical links at the lowest level of the PNNI hierarchy and logical links at all other levels. Without loss of generality, we denote the link weight as the link-delay metric, and use link delay as the focus metric of our algorithm.

Topology-Clustering Algorithm

The network-clustering algorithm is used to partition the network into a hierarchy of sub-domains and to construct peer groups. The recursive nature of the PNNI hierarchy dictates the algorithms to use recursion to partition the network into levels of peer groups in a bottom-up fashion, in which process peer groups are abstracted into logical group nodes.

It usually starts with a delay matrix at the lowest PNNI level, i.e., the physical switches and links as the nodes and edges in the graph G (V, E, Ø) with total of N nodes. Because this matrix covers the entire network, there are no border nodes to have links to the outside of the network. Matrix elements are initialized as $c_{ii} := 0$ and $c_{ij} := d_{ij}$ where $i \neq j$ and $0 \leq i, j \leq N$, and d_{ij} is the direct link-delay between nodes i and j. If there is not any direct link between node i and j, set $c_{ii} := \infty$.

$$C^0 = \begin{pmatrix} 0 & c_{1,2} \ldots & c_{1,j} & \ldots & c_{1,n} \\ c_{2,1} & 0 \ldots & c_{2,j} & \ldots & c_{2,n} \\ \cdot & & & & \cdot \\ \cdot & & & & \cdot \\ c_{i,1} & c_{i,2} \ldots & 0 & \ldots & c_{i,n} \\ \cdot & & & & \cdot \\ \cdot & & & & \cdot \\ c_{n,1} & c_{n,2} \ldots & c_{n,j} & \ldots & 0 \end{pmatrix} \Rightarrow C^1 = \begin{pmatrix} pg^1_1 & c^1_{1,2} \ldots & c^1_{1,j} & \ldots & c^1_{1,m} \\ c^1_{2,1} & pg^1_2 \ldots & c^1_{2,j} & \ldots & c^1_{2,m} \\ \cdot & & & & \cdot \\ \cdot & & & & \cdot \\ c^1_{i,1} & c^1_{i,2} \ldots & c^1_{i,j} & \ldots & c^1_{i,m} \\ \cdot & & & & \cdot \\ \cdot & & & & \cdot \\ c^1_{m,1} & c^1_{m,2} \ldots & c^1_{m,j} & \ldots & pg^1_m \end{pmatrix}$$

Next, an affinity function is defined, usually a function of hierarchical levels, that cluster logical nodes together into peer groups. The affinity criterion can be a number to limit the maximum number of logical nodes in a peer group, or it can be a function of the QoS parameter of concern. We will use link delay in our case. The link can be either physical or logical. A modified version of the Kruskal Minimum Spanning Tree has been used to construct a PNNI hierarchy of peer groups, i.e., a minimum spanning-tree forest as a collection of peer groups. Once the partitioning is done at one hierarchical level, the PNNI complex node transformation algorithm is invoked to convert the peer groups into the PNNI Complex Node Representation. Then we can replace the matrix blocks—entries of the same peer group—with a single delay element, as the post-aggregation matrix C^1 shows. Generally, for peer group i in C^1, entry c^1_{ii} will be the default spoke value of the PNNI Complex Node Representation of pg^1_i, and c^1_{ij}, the delay between peer group i and peer group j at the same PNNI hierarchy levels. For the nodes not belonging to any of the spanning trees, we can leave them to be included in peer groups of a higher PNNI hierarchical level later. This clustering process continues to recursively aggregate and build higher-level delay matrices C^2, C^3, …, until the whole network is contained in one peer group, the highest level C^D.

Based on the affinity criterion used in the partitioning process, link delays among nodes of higher-level groups are greater than those among of lower-level group nodes. In other words, at a certain hierarchical level, peer groups X, Y, and Z are viewed as delay matrix blocks, and delay matrix block A, i.e., a higher-level peer group, $a_{ij} \geq \max\{x_{mn}, y_{pq}, z_{rs}\}$, where $a_{ij} \in A$, $1 \leq i, j \leq |A|$, $x_{mn} \in X$, $1 \leq m, n \leq |X|$, $y_{pq} \in Y$, $1 \leq p, q \leq |Y|$, $z_{rs} \in Z$, and $1 \leq r, s \leq |Z|$, $|A|$, $|X|$, $|Y|$, and $|Z|$ are the sizes for the matrix blocks A, X, Y, and Z, respectively. Also, the sum of the intra-group bypasses is less than the inter-group cost, i.e., $\max\{b_{mn} \mid \forall\, m, n \in X\} < a_{mk}$, where $e(m, k) \in E$ is a valid link, $k \in A$, but $k \notin X$ and $m, n \in X$. In case this assumption does not hold at a certain level inside the hierarchy, make one group with larger b_{mn} into higher levels directly, i.e., a PNNI induced uplink, until the affinity function returns true for the new level or into the highest level of the hierarchy, where inter-group delay will no longer be of concern.

When we start the partition-conversion process, the inter-peer-group links are aggregated so that there is only up to one link between any two peer groups at a higher level. This is generally applied to any inside levels of the hierarchy. The only exception to this occurs at the lowest level, i.e., the physical level, where multiple links may exist among different peer groups as no aggregation is done.

The PNNI complex node transformation algorithm in the next section has been used to convert peer groups into PNNI Complex Node Representations. The algorithm recursively works on the delay matrix obtained from the aforementioned PNNI partitioning algorithm. It recursively partitions the matrix into disjoint sub-matrix blocks, aggregates each block into one logical node, and generate summarized matrices.

Due to the overlaps of the links aggregated, the delays of the spokes cannot represent the true minimum delay parameters from border nodes to the virtual nucleus. Bypass links are introduced to handle these exception linkages. Define an internal delay variable, equivalent to intra-node transit delay or internal queueing, so that a shortest path or minimum link delay can be used between two peer groups. Besides considering the cumulating of inter-group delay on the path, intra-group delays need to be included by each group on the path and to add to the path's delay when path is being built up.

$$d_{ij} = pg_i + d_{ij} + pg_j$$
$$d_{ii} = pg_i$$

With the topology-state database, an ATM CAC can quickly admit or reject a new network connection request based on the end-to-end delay QoS requirement of the new connection request and path delay information in the local topology database. The other QoS metrics can replace link delays previously used to enable other end-to-end QoS requirements of new connections.

Topology Aggregation

In the following sections, we present a new approach for finding the logical links of the star networks [2]. This involves the algorithms that aggregate a graph G = (V, E, B) into a star network with bypasses, where V is the set of nodes, $B \subseteq V$ is the set of border nodes, and E is a set of bidirectional links among the nodes in V.

Different aggregation schemes have been proposed for the ATM Forum's PNNI hierarchy, such as star, minimum spanning tree, and random spanning tree. Based on the simulation study, the star-based aggregation schemes have consistently performed better than schemes based on minimum spanning tree and random spanning tree [7].

The first step is based on Dijkstra's shortest path algorithm. It has positive edge weights and calculates the shortest paths among border nodes in an undirected graph. For a given peer-group representation in a graph G represented with adjacent lists, the purpose of this first step is to find the shortest paths from each border node to all the other border nodes. First, transform the given graph into a complete mesh, then search for the shortest path for each border node as follows: From each source border node, initialize a weight function of each node to infinity, which records the running delay from the single source to other nodes. Then relaxation technique is applied to obtain the shortest paths. These minimum delay results are stored in a delay matrix. At the end of the step, the delay matrix will hold the minimum end-to-end delays among all border nodes.

The next step is to transform the delay matrix into the PNNI Complex Node Representation—i.e., a star representation with bypasses—and store it in an enhanced delay matrix. Due to the overlaps of the links aggregated, the delays of the spokes usually cannot represent the true minimum delay parameters from border nodes to the virtual nucleus. Bypass links are introduced to handle these exception linkages. The following steps will generate the bypasses to make exception linkages available to the routing algorithms.

The process starts with a node set with one border node, then repeat the follow step until all border nodes are in the set: Find the longest delay link that connects nodes inside the set and outside. If the link delay is less than the existing spoke of the inside node times two, reassign this spoke to the half of the cross-link delay and add the outside node into the set. After the spoke is calculated, validate all exception bypasses into the enhanced matrix; if there is a link between two border nodes having a delay significantly less than the traversal delay through the nucleus, add this link as a bypass to the matrix positioned by the node IDs of the two nodes.

At the end of these processes, we have successfully aggregated the peer group with arbitrary topology into the ATM PNNI Complex Node Representation.

Conclusion

With these clustering and transformation algorithms, network-topology–state information can be aggregated efficiently to build the ATM PNNI hierarchy. With the routing tables built with these aggregated QoS parameters, an ATM connection admission control can quickly admit or reject new network connection requests based on the end-to-end QoS requirement of the new requests by simply performing some routing-table look-ups.

References

1. ATM Forum Technical Committee, "Private Network-Network Interface Specification," Version 1.0 (PNNI 1.0) (March 1996).
2. Y. Tian, I. Chu, and H.A. Chen, "Topology Aggregation with End-to-End Delay Guarantee in ATM PNNI Routing," in Proceedings of the 16th ICAST 2000 (Batavia, Ill.: 2000), 97–101.
3. I. Iliadis, "Optimal PNNI Complex Node Representations for Restrictive Costs and Minimal Path Computation Time," IEEE/ACM Transaction on Networking, vol. 8, no. 4 (August 2000): 493–506.
4. A. Iwata and H. Suzuki, "QoS Aggregation Algorithms in Hierarchical ATM Networks," in Proceedings of IEEE ICC '98 (Atlanta, Ga.: June 1998), 243–248.
5. W.C. Lee, "Minimum Equivalent Subspanner Algorithms for Topology Aggregation in ATM Networks," in Proceedings of IEEE '99 (1999), 351–359.
6. I. Iliadis, "PNNI Clustering Under Node Mobility in ATM Networks" (1998).
7. J. Liu, Z. Niu, and J. Zheng, "Two Optimization Approaches for Parameter Determination of PNNI Complex Node Model," in Proceedings of IEEE GLOBECOM '99 (1999), 1,066–1,070.
8. B. Awerbuch and Y. Shavitt, "Topology Aggregation for Directed Graphs," IEEE/ACM Transaction on Networking, vol. 9, no. 1 (February 2001): 82–90.
9. I. Iliadis and P. Scotton, "Transition Matrix Generation for Complex Node Representation," ATM Workshop (1999).
10. E. Felstaine and R. Cohen, "On the Distribution of Routing Computation in Hierarchical ATM Networks," IEEE/ACM Transaction on Networking, vol. 7, no. 6 (December 1999): 906–916.
11. R.K. Ahuja, T.L. Magnanti, J.B. Orlin, *Networks Flows*. (Englewood Cliffs, N.J.: Prentice Hall, 1993).

12. L. Georgiadis, R. Guérin, and V. Peris, "Efficient Network QoS Provisioning Based on per Node Traffic Shaping," in Proceedings of IEEE INFOCOM '96 (San Francisco, Calif.: March 1996), 102–110.
13. E. Felstaine and R. Cohen, "On the Distribution of Routing Computation in Hierarchical ATM Networks," IEEE/ACM Transactions on Networking, vol. 7 (December 1999): 906–916.
14. C. Huitema, *Routing in the Internet*. (Englewood Cliffs, N.J.: Prentice Hall, 1995).
15. B. Awerbuch, Y. Du, B. Khan, and Y. Shavitt, "Routing through Networks with Hierarchical Topology Aggregation," in Proceedings of IEEE ISCC '98 (Athens, Greece: 1998).
16. Y. Tian, I. Chu, "ATM PNNI Clustering for Topology Metrics," ISCA 14th International Conference on Computer Applications in Industry and Engineering," CAINE-2001 (Las Vegas, Nev.: November 2001).
17. S. Chen, "Routing Support for Providing Guaranteed End-to-End Quality-of-Service," (Ph.D. thesis, University of Illinois at Urbana-Champaign, 1999), http://cairo.cs.uiuc.edu/papers/Scthesis.ps.

Service Quality in Large-Scale DSL Deployments

Didier Cop
Vice President, Marketing
Amplify.net

Service Quality and Assurance

What are today's requirements for the digital subscriber line (DSL) infrastructure to deliver service quality and assurance? This article will discuss this question, as well as technology requirements for the access network.

Service-level agreements (SLA) receive a great deal of attention in this industry. But SLAs are a myth. They are a myth because they only work if there are teeth in them, a fact that that really applies to the corporate marketplace, with huge contracts and true SLAs. DSL is a residential/small-business play, and has so far been a best-effort service offering. As we will see, if implemented correctly, service quality is a good substitute for SLAs.

Quality of Service

So what should really be service quality? Taken from a global perspective, an SLA is mostly understood as a written contract wherein particular performance criteria are set out. But when considering the current DSL market, talking about an SLA really means being concerned with meeting the satisfaction of the subscriber and the users at this subscriber. A distinction can be made between the two. The subscriber is the decision maker in any particular household, while users are the people using the service and can be very influential on the subscriber. Many subscribers and users have been recently disconnected, initially because of a failing Internet service provider (ISP) or a failing data local-exchange carrier (DLEC). It is increasingly important that DSL users be happy with the quality of the service they receive. They will then spread positive messages, rather than negative ones.

Quality of service also means meeting the internal metrics as specified by the service provider. For example, if a subscriber might have signed up for a 640 kilobit (k) service, the internal service-provider specifications could be that the user should never be below 128k but get 1 megabit per second (Mbps) at off-peak times. This is an informal SLA, wherein the marketing folks decide what the service quality should be; it is not on a contractual commitment, but it is one way to define the service level.

DSL is a service crossing several boundaries, from the end customer to the final destination on the Internet. It goes through several networks such as the access network or the metro network, and real SLAs here are required. We are talking about bigger monies and bigger contracts. The ISP is actually offering service through an open-access environment and shares the same access infrastructure with several other ISPs. ISPs are dependant on the services metro providers offer.

Subscriber Satisfaction

Does the subscriber get what he or she pays for? For standard Internet access, what is the phenomenon observed by the subscriber or the user when the contract is not delivered? In most cases, he or she will not notice unless a software performance tool is used because the service is data, not real-time. The same goes for services around security. The only time when the user will feel the contract is being breached is if an intrusion occurs, and while the risk is there, the event is infrequent.

Today's new communication value-added services—like voice and video—are being introduced among well-established metrics of quality such as toll-voice and broadcast television and DVD. Here the SLA contract is very simple—will the user get the quality experience he or she expects every time? That's what quality of service is about. In these instances, because of the user's well-established quality expectations, it is possible for him or her to feel immediately that the service contract is broken.

Service-Provider Specifications

For these real-time services, service providers should provide a high level of service quality—or no service at all. Historically, it is what customers have come to expect from regular phone service or the DVDs they rent. Sub-quality service will not be good enough to keep the customer. Following are a few example guidelines for service specifications:

- Bronze customers should never get less than 20 percent of the service level they expect (i.e., 640k to 128k)
- Silver customers should never get less than 40 percent of the service level they expect (i.e., 640k to 256k)

- Voice should be 64k for business (worst case), 32k for residential, and denied if not possible
- Video on demand should be a DVD–like experience

Infrastructure Requirements for Paying Value-Added Services

The first infrastructure requirement must be the ability to guarantee that the actual level of quality is being delivered for each of the services, each time. The second is maintaining the quality end-to-end—from the end-node at the subscriber to its destination, which can be the ISP or a video server or a voice gateway, which is not controlled by the ISP, but rather is controlled by the access provider or another partner. Access providers have a further challenge as they are serving or will shortly serve several ISPs and application service providers (ASP) through the same infrastructure. They will have to deliver SLAs to each that enable them to differentiate and compete in the market.

Technology Considerations

Quality of Service in the Access Network

In *Figure 1* a distinction is made between the last mile, which is everything from the customer to the access device in the central office or in a digital loop carrier (DLC), and the metro or long haul, which links all the subscribers to various ISPs.

Figure 2 shows quality-of-service technologies, used in asynchronous transfer mode (ATM) or Internet protocol (IP) networks, such as multiprotocol label switching (MPLS) or resource reservation protocol (RSVP). None of them have been engineered for the last mile, where bandwidth challenges dominate. It is rather easy to scale the metro network, but it is very difficult to progressively add bandwidth in the last mile.

Last-Mile Aggregation: Internet Access

Historically, broadband has simply been Internet access, with a very simple playing field—users were connected through modems to a digital subscriber line access multiplexer (DSLAM) before reaching the metro network. Now we have to face adding multiple real-time services serving multiple ISPs. For DSL, while it is now dominantly self-serving to the incumbent local exchange carriers' (ILEC) own Internet entity, ISPs still have a lot going for them and have an integral part to play in marketing, selling, and delivering services to the end customer. Open access is also emerging in the cable environment, where multiple ISPs are starting to gain access to the cable infrastructure owned by the cable companies (see *Figure 3*).

Multiservice Open Access

Access providers must be prepared to aggregate traffic between multiple points to and from the same subscriber. It will not be long before a subscriber will want to access several service providers for several types of services. Different users at the subscriber home or business will have different inclinations from the same PC and will want to use a plethora of services that all require different behaviors, all behind the same IP address.

This poses a formidable challenge to access providers wherein the first location just after the customer premises becomes the only real service aggregation point before it gets distributed to different destinations. This point that requires a lot of granularity in terms of quality of service. When an access provider is serving multiple ISPs, it will need to enable service differentiation and meet different service models and different quality-of-service requirements. *Figure 4* shows a multiservice open access environment with subscribers using multiple services.

FIGURE 1

End-To-End

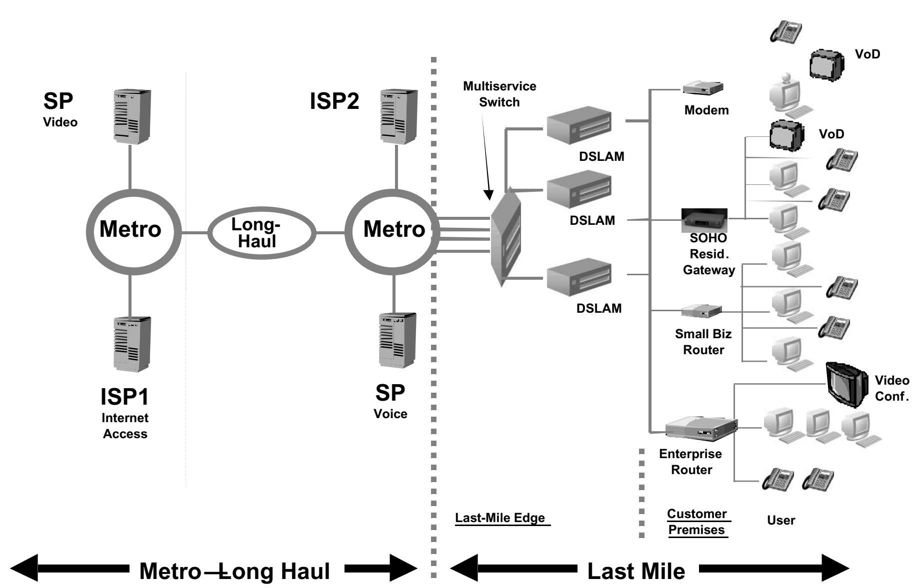

FIGURE 2

End-To-End Service Quality

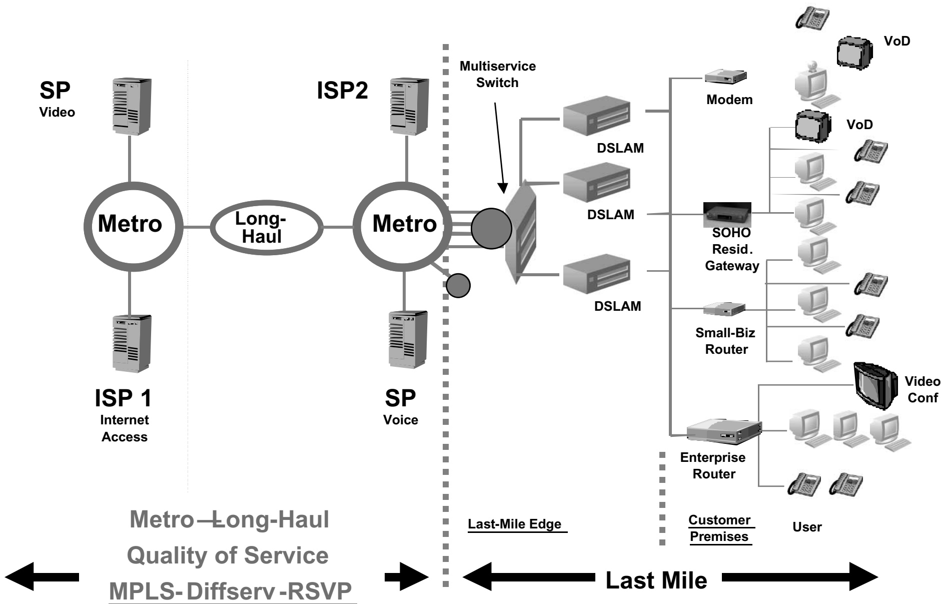

Each of the DSLAMs before aggregation to the multiservice switch is a potential physical congestion point (see *Figure 5*). Because the access infrastructure is going to be serving several ISPs, there will be a number of virtual congestion points, such as the total bandwidth allocated to each ISP and the classes of services that each ISP has established. It is likely that each ISP will have its own strategy of what a silver service should be and how much over-subscription it wants to implement.

When examined from the user's PC perspective—that is, looking at the path that the packet flow is going to take—this is all a very complex situation. For all real-time or near-real time services, the correct amount of bandwidth neces-

FIGURE 3

Internet Access

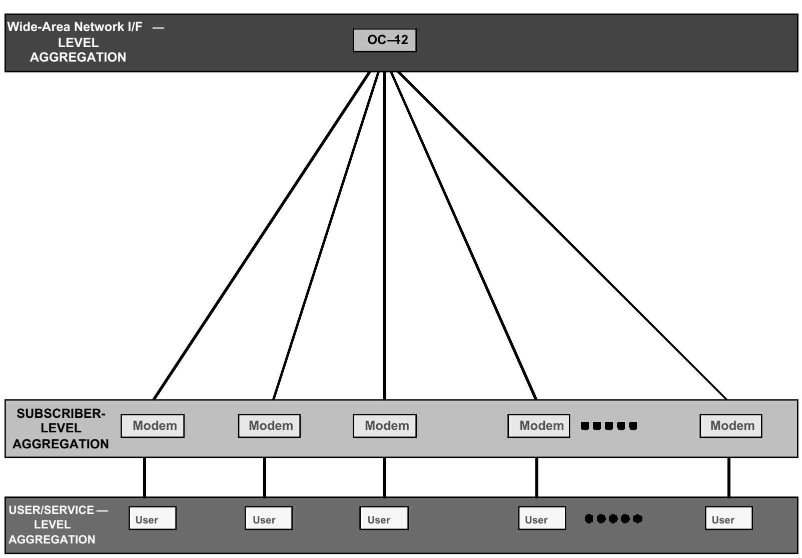

FIGURE 4

Multiservice

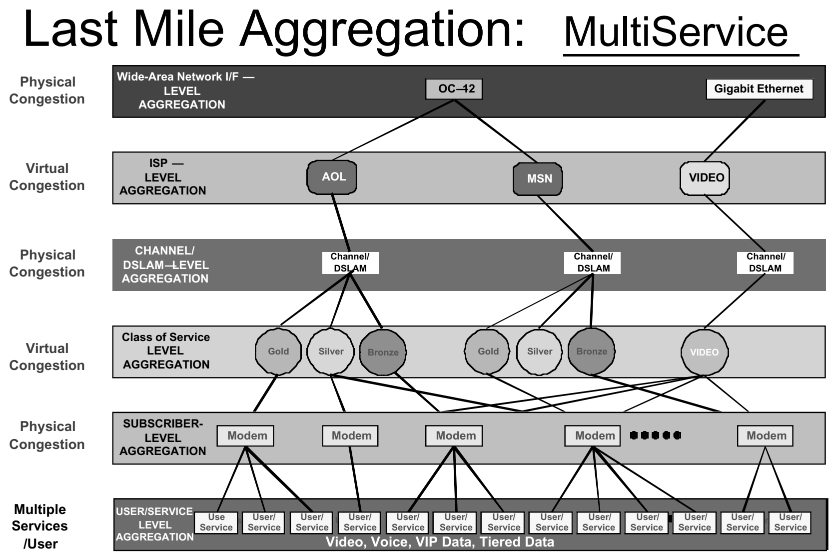

sary to meet each of the services' quality definition must be guaranteed, from the PC through the last-mile infrastructure, through each of the congestion points, virtual or physical. However, this deterministic guarantee needs to be implemented in a way that makes bandwidth available at all times for non-real-time traffic if not used.

Another important concept is fair bandwidth attrition for non-real-time traffic—in other words dealing with what happens to best-effort data while all high-priority traffic is being served. As more services that consume bandwidth are introduced, data services are going to degrade if not enough attention is paid to them. Again, bandwidth in the last mile is not elastic. Access providers should use IP traffic-management technology to make sure that each of the subscribers receives a fair amount of bandwidth.

FIGURE 5

The Multiservice Switch—A Potential Physical Congestion Point

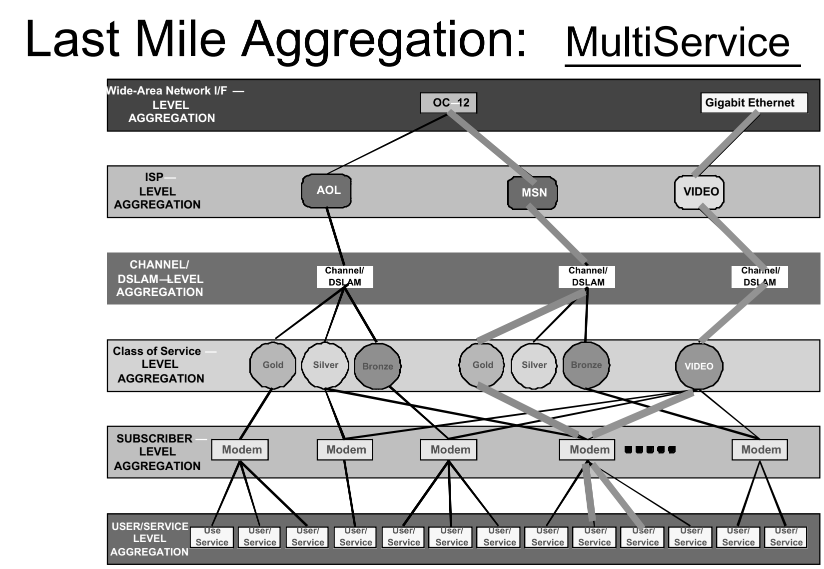

Multivendor OSSs: The Challenge Is Speed

George Cray
Executive Director
Telcordia Technologies

Topics

This paper considers the problem of multivendor operations support systems (OSS) and multivendor networks from the perspectives of a network operator, a network equipment vendor, and an OSS provider. Multivendor networks present real challenges, and foremost among them is to deliver a full solution, especially from a customer's point of view. To examine some of the demands, it is best to consider the different perspectives around multivendor networks.

The paper also examines what is needed to provide OSS and network element (NE) interoperability—in other words, what is needed to make them work together. While many boast of "plug and play," the reality is that plug and play is never truly attained. There are techniques and standards that can shorten the time frame, but the integration effort remains formidable. A real example, one that involves digital subscriber line (DSL) technology, will provide some lessons learned regarding today's environment.

A Network Operator's View

Basically, network operators want everything. They require OSS solutions that support multiple services, voice over DSL (VoDSL) or voice over Internet protocol (VoIP), data, and Internet protocol (IP) virtual private networks (VPN). Ultimately they want to use networks to transport any service they sell, so trying to get to a truly multiple-technology, multiservice network is the goal. Multiple technologies clearly come into play. These include traditional time division multiplex (TDM) technologies, synchronous optical network (SONET)/synchronous digital hierarchy (SDH), optical-based technologies, frame relay (FR), cell relay going across that, IP services, and asynchronous transfer mode (ATM). Operators ultimately are involved with multiple-vendor networks. That could be because they want different options from a purchasing perspective. It could mean that they acquire different operators' networks along the way.

Network operators do not want to be limited. While they go after solutions that handle multiple services, multiple technologies, and multiple vendors, they need to be scalable and to deal with transfer of assets very cleanly. Of course, they want it easy to use, easy to maintain, stable, and affordable. While the targets of ease of use, ease of maintenance, and stability might be attainable, they might make the solution unaffordable. It really is a challenge to try to make these networks operate the way a network operator would like them to.

The Network-Equipment Vendor

Network-equipment vendors want to be first to market with new technology. No one can blame them for trying to roll out new products first. After all, is it better to be first or second in a market? What that means from an integration perspective is that, typically, a lot of these new technologies or even the service offerings are ahead of any standards, and network-equipment vendors are actually hoping to set them. When it comes to new technologies, standards are few and far between.

Equipment vendors also want to differentiate their product from the competition. So, again, even if a set of standards exists, certain companies are trying to push the envelope and create new features and functions that maybe are not addressed in the standards. Maybe vendors address them differently from what is described in the standards. Equipment vendors want to innovate constantly. Thus, even if the standards start falling in line with the products rolling off, situations will invariably arise where someone will come out with that next new capability, and many new releases will follow. From a network-management point of view, element management systems (EMS) are quickly trying to churn out new releases and new interface specifications, not to mention acquisitions of new products.

For example, network-equipment vendors might develop strategies that try to bring innovations all under one EMS, whereas, previously, a number of NEs may have been managed by completely separate EMSs. However, when one company aggressively pursues market-share, that company, out there acquiring other companies, will have different strategies for supporting the elements underneath EMSs. Vendors want to stay in front, and they do not want to wait for OSSs to catch up. They do not want to wait for operators. They are being pushed by operators to give the next greatest capability.

An OSS Provider's Opportunity

For an OSS provider, one of the big challenges is that, initially, many of the network operators may be doing things with hardly any OSSs. They may be trying to put in equipment, maybe in a stand-alone mode segregating their network. They may be in a situation where they are in low growth at a certain point, pre-provisioning or manually activating the network itself. But as they grow they want to become fully automated, dynamically provisioned, and large-scale, with multiple services that will create a high-growth environment. That is why they are going after the OSSs.

The problem that OSS providers need to account for is that operators are already operating. So, OSS providers cannot impact the current ability to provide and maintain service. They must be able to provide automated data conversions to the new OSS. Data may be in another OSS or, more likely, in a spreadsheet system. The trick is figuring out how to get that data into the new OSS. The network operator will not own the entire network. There will be many handoffs with other trading partners, which only complicate trying to provide that end-to-end service.

Hence, vendor selections and service offerings will change, and they typically change on or before initial deployment of OSS, which is yet another challenge. What is stipulated in a contract for delivery versus what is ultimately put into service typically does change, even if it is in a three- or six-month window. That is simply the reality of what it takes to win business today. Everyone has a different perspective, yet the challenge is to make it all work three months from the start of signing a contract or in *x* number of weeks after a decision is made to purchase a piece of equipment.

OSS–NE Interoperability

A number of issues need consideration regarding OSS and NE interoperability. Integrating OSSs with NEs to simplify network operations requires communications protocol standards (for example, X.25, transmission control protocol [TCP]/IP, and open systems interconnection [OSI] Layer 7). Obviously, the idea is to synchronize with what the protocol is going to be. The problem is that picking a standard does not mean that it will work automatically. Unfortunately, these protocol standards allow for implementation variations via parameters and protocol data. Simply selecting OSI Layer 7 or TCP/IP does not mean that all standards are equal.

Similarly, OSS and NE interoperability requires message protocol standards to define the syntax, message structure, and punctuation. Again, vendors usually adhere to message protocol standards, but minor deviations exist. A good example involves common object request broker architecture (CORBA), with such variants as VisiBroker® and Orbix®. Likewise, there are different versions of simple network-management protocol (SNMP). While message protocol standards will help interoperability, they are not a guarantee that systems, whether a network-management system (NMS) or an EMS, will really plug together.

Interoperability also depends on data and information-modeling standards to define how to interpret message contents. But several management information bases (MIB) exist, both with SNMP and common management information service element (CMISE). CORBA has an interface definition language (IDL). Certain types of interfaces can be standardized, but all standards are subject to interpretation, and all allow for extensions. An illustrative example involves ATM where SNMP was being used. The SNMP standard for ATM called for about six traps from a surveillance perspective, and the ATM switch in question had more than 100 traps. A good portion of those the network operator really might want to use, and it does not mean all must be used, but care should go into determining what will really be used and what needs to be standardized for that network operator. Simply having a standard will not allow network operators to deliver all the functionality unless they get beyond that standard and talk about the underlying EMS or NE and determine what is going to be used by the network operator to take advantage of some of the features and capabilities that were in that NE.

Integrating OSS with NEs also entails equipment and network modeling, dealing with physical to logical representation of equipment and networks. Typically, NEs are fundamentally designed and engineered differently. Some things are relatively standard, such as physical layout (bay, shelf, slots, and ports), features (types and quantities of drops), and certain supported architectures (chain, ring, or hybrid), but every vendor chooses to implement things slightly differently. That must all be accounted for from both a physical and a logical perspective to do assignments and deal with alarms.

Interoperability testing (verification according to specifications) is another integration issue concerning OSS and NEs. What stands out foremost is the challenge of speed. Network architects spend much of their time actually defining the requirements and doing the testing. Time is not spent writing the code or putting the interfaces together. Getting through the actual requirements—what the interface will be, what it will be standardized on, and what the extensions are—takes the most time. Once the code is written, testing must follow, and testing really is where the rubber meets the road. That is what determines whether the integration will work. Testing also uncovers numerous other things not obvious from reading the documentation. It is clear that documenting these northbound interfaces is not the first thing on an equipment vendor's list of things to do, and, therefore, trying to keep up with the capabilities in the EMS or the NE becomes a challenge when looking solely at documentation. For interoperability testing, it will be necessary to verify the synchronization of logical representation in support of multiple–OSS flow-through environments.

Ultimately, the end-to-end solution will deal with some methods and procedures. It is critical to establish how to handle irreconcilable differences. There will be out-of-sync conditions, and, through order discovery, network operators will find things that are not the way they thought they were. In these different areas, methods and procedures may come into play, and to think differently is foolish.

A Real Example

A real example will provide an overview of a network that was in place when management efforts began (see *Figure 1*).

FIGURE 1

DLEC DSL OSS Solution

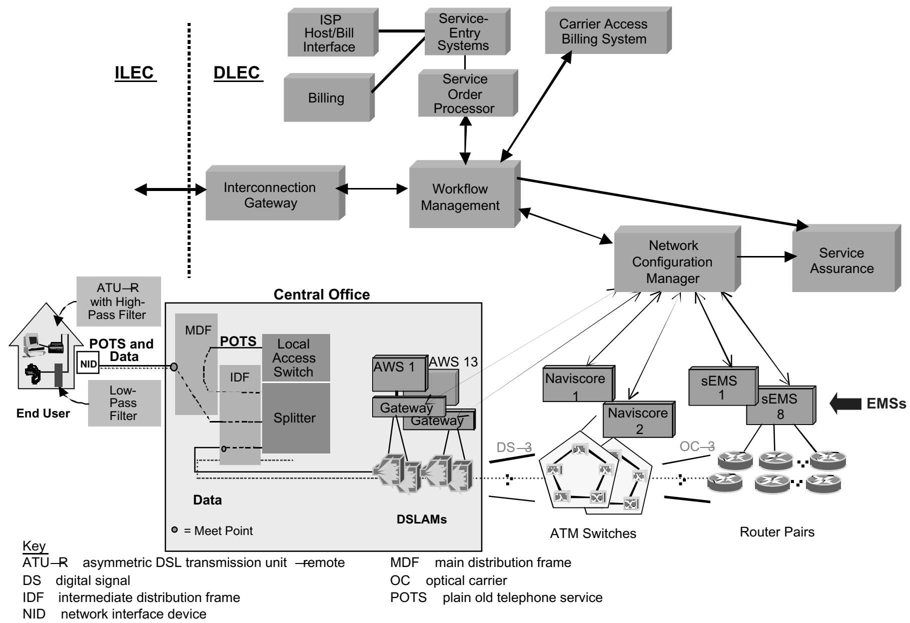

In this case, a data local-exchange carrier (DLEC) wanted a DSL OSS solution. The focus in this example will be on the network configuration manager, which is basically an NMS for configuration or provisioning. It talks to a multivendor network and to EMSs, all under the guise of doing DSL provisioning. The DLEC had NEs that were performing the DSL access multiplexer (DSLAM) capability—in this case, an Alcatel 1000 ATM subscriber access multiplexer (ASAM). There were also RedBack routers and a Lucent ATM switch and that were managed through a Nortel EMS.

The network operator was looking to turn up service in a flow-through environment. There is no need to concentrate on the rest of the OSSs involved because the goal here was to activate that network. What must be done to provision in a multivendor environment?

DSL Activation: The "Real" Process

Figure 2 breaks down the configuration process. When actually messaging to the EMSs, the interaction is very different based on whether it is the DSLAM, the ATM switch, or the router.

The first step is a request for network activation. The "–1 Day" refers to a day before the due date, which can be set arbitrarily. In this case the network operator chose to activate a day before the due date, and the orders were sent down to the network configuration manager, which is really only an NMS dealing with the provisioning aspect across a multivendor environment. That network configuration manager can now talk to the various EMSs through a variety of protocols: transaction language 1 (TL1), SNMP, and CORBA. Recall that this is in a supposedly homogeneous network, but that different standards apply.

In the case of talking to the DSLAM EMS, times have been greatly improved. The 1000 ASAM could basically handle the orders only in series across DSLAMs in parallel, resulting in long wait times to set up the cross-connect (XC). Then timers were set. If no response was received, another attempt was made. Often this happened several times, averaging a three-minute total lapse time to get, potentially, a successful activation. Now this typically takes less than a minute, and the typical XCs are done in eight to 15 seconds. After an EMS is placed in an automated flow environment, the performance becomes a serious factor, and operators start asking what this does to their overall ability to provision services. Initially, the interface was serial, meaning there were great limitations on the number of orders it could process. Problems with timeouts required looking at that interface. The network operator, the equipment vendor, and the OSS provider worked to reduce that into a much more reasonable time frame.

The ATM switch had a different set of issues. There the northbound interface dealt with parallel processing, meaning many orders could be sent down through the EMS. In that case there were different timeouts, which were set up based on the capability of that EMS to process orders. A request for a permanent virtual circuit (PVC) operated on a 4.5-minute lapse: 30 seconds for the first try, with four 30-second retries each with 30-second waits between them. At the router EMS, orders were processed serially, which was a problem, and the timeouts to set up the XC were set to 20 minutes because the goal was to provision all three of these components successfully.

If that was not done, the network configuration manager sends messages to each EMS to tear down the circuit (rollback) and a request for manual assistance (RMA) is sent

FIGURE 2

DSL Activation: The "Real" Process

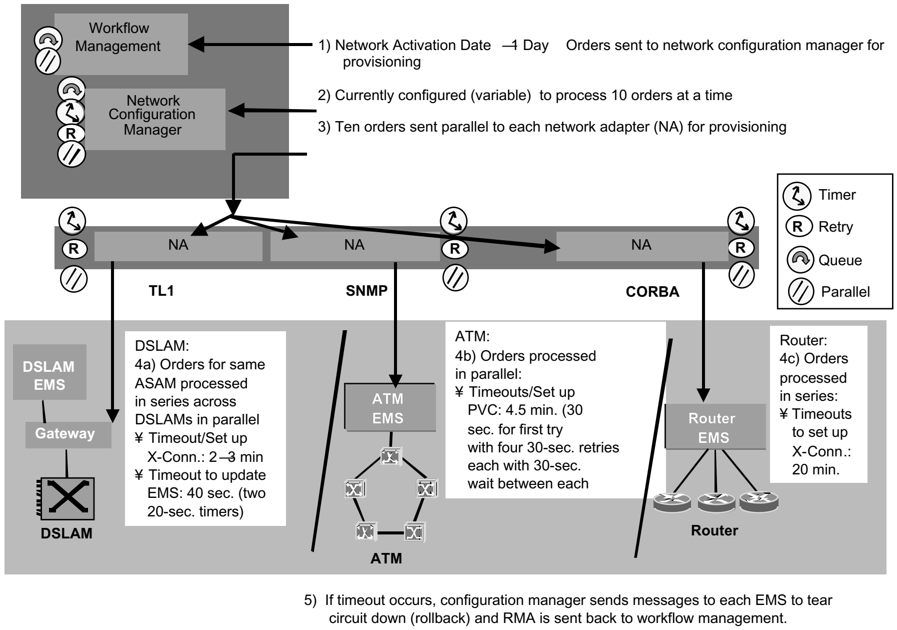

back to workflow management. If the DSLAM or some portion of the ATM network was activated, but not all the way through to the router to the Internet service provider (ISP), ultimately all that was provisioned has to be rolled back and manual assistance has to be requested. Network operators want these services to be provisioned end to end and do not want to encounter problems with the router or the DSLAM. Once everything is rolled back, the network operator can choose to have the systems retry or to have a user provision manually.

Lessons Learned

Two of the lessons learned involve network health and the OSS software solution. Network health depends on capacity planning and scalability to ensure that the network and the EMSs were able to provision the expected volume. This lesson is often learned in production and may not be known ahead of time. EMSs tend to be job aids for engineers and technicians, if initially no thought is given to northbound interfaces. In that case, the expectation of that EMS as a job aid versus an integral part of flow-through provisioning is very different. The idea of supporting simultaneous (versus serialized) requests is critical. Scalability depends on the ability of the number of NEs actually supported underneath a given EMS.

Obviously stability and availability are important. Both the NEs and the EMSs must be stable and highly available for provisioning with service assurance actively monitoring the EMSs and NEs. If EMSs go off-line for significant chunks of time and are not able to be queried or given provisioning orders, trying to provision in a flow-through environment will be difficult. That lesson is something probably learned in a lab or through testing rather than through reading any documentation.

The OSS software solution needs solution scalability. Clearly in the real example above, the operator wanted to deal with all order types and all order passes. There were situations with local service requests (LSR) and access service requests (ASR) with multiple providers providing pieces of the network that got in the way of providing an end-to-end service. Solution stability must deal with error conditions, and it must involve a process to handle rollbacks and to keep volumes down. Solution operability must contend with network and EMS unavailability and timeouts, and, hence, solutions should be engineered to limit the amount of error conditions. It must provide an increased view of network data, a single view of an order across the total environment, and message queuing (user interactions versus flow-through). That is why network operators are buying an automated set of systems, so that the network can be provisioned error free. Having the right interface specification is not the whole answer, but improved OSS is tied to how the NEs and EMSs behave. Availability is the heart of the matter.

An OSS Answer for Increased Speed

One of the answers to how OSSs will improve NEs and EMSs is a fully integrated network database. The value is in single data entry and single data retrieval. Obviously, using the network as the database will be helpful, but the trick is in crossing functionality—whether it is provisioning, assurance, or performance—and in not relying on separate databases. The goal is to have a single network touch point for all EMS interactions focusing on available standards. This will simplify network interaction. But if the goal is touching the network once, a common platform is the best path to take. A common platform will have a consistent standard operating environment (SOE), a common graphical user interface (GUI), system interface specification and documentation, integrated capacity planning and performance, and improved integration. Fewer servers will be needed, systems administration will become easier, and user interconnection will be simplified, multifunctional, and ready for system integration. A fully integrated network database, a single network touch point, and a common platform will make implementation more rapid and hopefully will simplify the lives of users and network operators. Replication requires simplification.

Summary

Multivendor networks are a reality, but their management is difficult. There are no quick steps to get around the interoperability testing to understand how EMSs behave. Following the element-management layer (EML) and network-management layer (NML) standards is helpful, but to think that they will solve the problems is naïve. Close relationships are needed between the OSS vendors and the equipment vendors to work on the expectations of what the EMS should or should not be doing. Following standards is not sufficient; the question of interface expectation needs to be addressed. Managing multivendor networks goes far beyond gaining connectivity to multiple vendors' NEs or EMSs.

Billing for Wholesale Services: Just Like Retail, or Is It?

Becky Dancy
Director, Business Development
Abiliti Solutions

Like a stadium used for both baseball and football, a billing and customer-care system used for both wholesale and retail billing can present challenges. While the stadium must change field configuration, sports equipment, and the team players, the software system must support differing business models, business drivers, and network architectures. A billing and customer-care system supporting both types of services must balance differing invoice formats, account structures, customer access, reports, taxation, usage guiding, collections capabilities, rate plans, reconciliation, and so on.

Customer demographics and expectations for wholesale and retail services vary greatly. For example, a wholesaler may have 40 to 50 large customers while a retailer's customer base can number in the hundreds of thousands. And, while the retailer's customers may charge only a hundred or so dollars each month, the wholesaler's customers can have thousands of dollars of monthly billing. The differing service-level expectations and functionality needs for retail and wholesale services make their mark across all components of a billing and customer-care system.

To illustrate, existing retailers entering the wholesale market may assume that they can simply use their traditional retail systems to support their new business. However, this assumption may not be accurate. While traditional retail systems may very well get the wholesale invoices out the door and enable the providers' customer-service representatives (CSRs) to answer simple customer inquiries, they may not efficiently or cost-effectively meet the providers' actual wholesale business objectives.

Product Management

All billing systems must be able to support product creation and maintenance. So, why would a wholesaler need any special functionality in this area? If a billing system can support local, long-distance, or data services in a retail market, it should have no problems supporting wholesale flavors, right?

While this idea might be correct in many situations, wholesale businesses often extend beyond traditional telecommunications services. Services such as 800-number translation or caller-name delivery service are not standard retail telecom products, and a traditional retail system may not be able to define or rate these services. Specifically, most traditional retail billing systems support only two-digit charge amounts, such as $4.95 for caller ID. Database services such as the ones mentioned may require rates up to four or five digits (i.e, $0.00175 per look-p). The wholesaler that maintains a caller-name database and offers that service to other telecom providers must have a billing system that can define and maintain this "unique" service.

Additionally, typical wholesalers use individual case basis (ICB) pricing for a large percentage of their customers. So, a wholesale billing system must be able to efficiently support special pricing for each customer. While a retailer probably places more importance on a system that supports global, mass pricing updates, a wholesaler needs a simple and quick way to make individual customer-rate updates. This is critical in many situations. For example, a wholesaler of international long-distance services with 10 percent of revenue from delivering traffic to Germany discovers that a competitor has introduced a lower rate to Germany. If the billing system is not able to support ICB updates quickly, the wholesaler is likely to lose a large portion of its business.

For a retail provider, the ability to offer the infamous "buy one, get one free" or a "purchase local and long-distance service and get Internet free" promotion is critical. This type of packaging is not really crucial for wholesalers. The ability to provide tiered and stepped volume and contract-based discounts, however, is essential. A wholesaler must be able to discount based on term of contract and committed usage (revenue, minutes of use, packets, etc.), and be able to discount not simply on the current month's activity, but also on the aggregate usage over a period of months or years.

Order Management and Provisioning

Billing and customer-care systems must be able to support order entry for all the services that they plan to offer, both traditional and nontraditional. This includes the ability to capture the data elements for each service type, as well as the ability to manage the ordering and provisioning process.

For instance, in a wholesale environment, where service providers are likely be involved in the final testing of the circuit or service, the billing system's order-management function must allow the capture of the end user's contact information, as well as the wholesaler's customer. It should also allow the addition of new provisioning steps to include end-user testing—without significant programming. Further, a wholesaler needs to be able to notify its customers—either via an application programming interface (API) or a manual process—that their service has been activated and tested, so they may begin billing their end users.

An increasing number of service providers are offering their end users remote access to account information through remote dial-in, direct connect, or the Internet. This allows the customers to query order status and even perform order entry themselves. When taking this route, service providers must ensure that their billing systems allow easy navigation around account data and that they contain adequate security measures to prevent inadvertent access to unauthorized information, or changes that could negatively affect service levels.

Event Processing

Because wholesalers may process 500 million event records or more per month, their system must be able to support greater throughput than the amount traditionally required by retail systems. But, throughput is not the only event processing-related question that must be considered.

Another, perhaps even more important, question is, "Can the system guide and rate the various types of event records that need to be processed each month?" If the system cannot guide the records, it is irrelevant how great the throughput of the system is. To illustrate, consider the situation of the before-named wholesaler that provides a caller-name service. These event records may be guided by point code or a customer ID, rather than "traditional" keys such as automatic number identification (ANI) or an 8xx number. In other words, the wholesaler may need a billing system with a flexible event-processing engine, not one that just has a high throughput capability.

Taxation also plays a major role in terms of event processing because, in most cases, taxes are not applicable in the wholesale arena. Not only does this eliminate the requirement for an event-processing system that can tax efficiently and correctly, it eliminates maintenance of tax tables and the creation of multiple tax log reports.

Customer Care

Wholesalers typically service their customers with dedicated account teams, while retailers usually have centralized customer call centers with hundreds of service representatives. These divergent approaches result in different demands on a customer-care system.

Because retail customer-care centers are often driven to effectively process a high number of calls with minimal call duration, the application must facilitate efficient navigation. This feature is not as critical in a wholesale environment where, generally, the interactions between the service providers and their customers are not necessarily as time sensitive. While it is obviously just as important to have all the necessary customer information available to the representative of a wholesale provider, it is usually acceptable if the navigation to get to that information is not as efficient.

Other issues related to customer care are account structure and scripting. In most cases, wholesale clients require less complex account hierarchies than retail business customers—wholesalers typically do not have to be able to represent their customers in two or three tiers with complex relationships. Likewise, the use of scripting to support customer-care dialogs and additional product sales is a more prominent issue for retail businesses.

The data-entry requirements for wholesale versus retail systems vary as well. In fact, a typical retail system may capture a high amount of extraneous information that is not necessary for a wholesaler's basic customer profile. For instance, a consumer system may have data fields for a driver's license number and a social security number, which are not applicable for a wholesale business. For both wholesale and retail applications, being able to hide or modify customer profile fields so they are more applicable to a service provider's business is a critical capability.

An area that is receiving increasing attention in both the wholesale and retail markets these days is customer self care. Service providers in both segments view the ability to allow customers to access account information and perform basic self-care functions—such as remote, real-time access for order entry and order status; invoice preview and payment; and trouble reporting—as a competitive advantage.

Credit and Accounts Receivable

On the financial side of the business, retail systems typically have more demanding requirements than wholesale systems. Wholesalers usually do not need automated credit checks and credit scoring to support their business, while these features are critical to a retail provider. Additionally, since most wholesalers use account teams to support their customers, they have less of a demand for sophisticated collections modules. However, the wholesaler's system will need to be able to create standard financial, aging, and collection reports.

Similarly, many billing and customer-care systems offer a wide variety of payment options, such as lockbox, Web, credit card, and bank draft. While this range of options is often very attractive for retailers, wholesale providers may not necessarily need such a big range, and they can end up paying for functionality that they do not really require if they aren't careful.

Invoicing and Reports

The invoice is the key mechanism that a service provider has to communicate with its customer base—in some cases it is the only way. For wholesale customers, however, the invoice and its associated reports are vital tools for managing their business. The invoice is their single device for reconciliation and network-cost management. For example, in a wholesale environment, the customer may be reselling services from a single source or from multiple providers. Whatever the case, the client has to manage and reconcile the amounts he is charged with the amounts he is billing to his end customers. The analysis involved in this mission-critical function can be very complex, and an effective wholesale billing system is one that enables the process for the wholesaler's customers.

Therefore, wholesale invoice information must allow the customer to easily and accurately compare invoice information to financial records. In most cases, this will mean providing the information electronically and allowing customers to freely manipulate the data and massage it into layouts that cleanly match their own records or billing-system outputs. For example, the billing system may produce reports (or invoices) that sort all services by product type. While this is certainly an acceptable, logical approach, the customer may record financial or billing information by customer or service ID. To simplify this customer's reconciliation and analysis process, an invoice format that allows the resorting of information by unique requirements is needed.

Another key item in the wholesale reconciliation process is allowing the customer to tie the amount that they are being charged for a given service or product to the amount that they are billing their end users. Or, more basically, enabling them to easily ensure they are actually billing an end user for all of the services and products for which they are being billed. Doing this requires the creation of a common identifier that can quickly and easily tie the services together. In some wholesale cases, this is more complicated than it might first appear. In the long-distance or local environments, for example, where the common identifier in most cases can be the 10-digit ANI, this process is relatively straightforward. However, it can become much more complicated when the identifier needs to relate features or other additional charges to the base service. For instance, wholesalers of local services may sell local features *à la carte,* while the customer retails the same features as part of packages. Reconciling the individual charges on the invoice or report to a package charge in their records may be complicated.

In the data world, where circuits are often combinations of various pieces, the reconciliation process can be very complex. For instance, an asynchronous transfer mode (ATM) circuit may consist of a local-loop circuit and 5 megabits per second (Mbps) user network interface (UNI) port in Denver, a local-loop circuit and 10 Mbps UNI port in Kansas City, and the 5 Mbps permanent virtual circuit (PVC) midlink. Not only may data providers have multiple ways to bill the service—each charge broken out on the invoice, combining charges based on node charges and PVC charges, flat rate versus usage based, etc.—but they may also take different approaches to identifying each piece. Specifically, does the port have a unique ID, or does it carry the same ID as the local loop?

In most cases, the circuit IDs provided to the wholesale customer will not be the circuit IDs that the customer will use to identify the circuit for their customer. They may simply change the lead letters or numbers (for example, ID ABC1234 becomes XYZ1234) or they may implement completely different number schemes (either by choice or required by their operations support system [OSS]). The ability to tie these network elements (NEs) and their charges from the information a wholesaler provides on an invoice or report to their records is critical for reconciliation. A wholesaler must ensure that its billing and customer-care system makes the reconciliation process as simple and streamlined as possible.

Conclusion

In the increasingly competitive telecommunications industry, the features and functionality of a service provider's billing and customer-care system can provide true competitive advantages. When selecting a billing and customer-care system to support your operation, there are many questions that you must ask and considerations to balance—regardless of whether you are a wholesaler, retailer, or both. When evaluating your options, be sure that you weigh the availability of particular system features against how important those capabilities actually are to your business. A single system solution for both wholesale and retail operations may support your business, but you may have to accept some limitations. In other cases, the preferred solution may be to purchase two separate systems to support the complete list of business drivers and objectives. In other words, if you are both a wholesaler and retailer, you will need to strike a balance between the meeting the drivers for each business, the capabilities that a single system can provide, the cost of modifying a single system to support both businesses, and the cost of purchasing and maintaining two systems.

In any situation, be sure you match the system capabilities to your business needs, to ensure that you are spending your money wisely and are investing in those functions that will bring you the best return on that investment. In other words, don't pay for functionality that you don't need and, moreover, you'll pay dearly if you need functionality that you don't have.

QoS Support in the DMA Architecture: Implementation and Performance Analysis

Sajal Das

Professor, Department of Computer Science and Engineering, and Director, Center for Research in Wireless Mobility and Networking (CReWMaN)
University of Texas at Arlington

Subir Das

Research Scientist, Wireless IP Research Lab
Telcordia Technologies, Inc.

Archan Misra

Research Staff Member, Pervasive Security and Networking Group
IBM Research

Abstract

We have recently proposed the intra-domain mobility management protocol (IDMP) as a part of the dynamic mobility architecture (DMA), a two-layer hierarchical solution for managing the movement of mobile nodes within a cellular access domain. Like most Internet protocol (IP)–layer mobility management protocols, the basic IDMP specifications aim only to ensure uninterrupted roaming connectivity. This paper describes how the DMA architecture uses enhancements to IDMP to provide quality of service (QoS) guarantees for mobile users within a cellular access domain. DMA's QoS framework uses a modified form of the Differentiated Services (DiffServ) architecture, with a centralized bandwidth broker (BB) performing admission control and resource provisioning for different traffic classes within the access domain. We present the enhancements to the signaling and functional architecture necessary to seamlessly migrate QoS guarantees as a mobile changes its point of attachment. After describing our Linux and FreeBSD-based prototype implementation of this framework, we present performance results obtained on our testbed.

I. Introduction

IP–based mobility management solutions, such as mobile IP [Perkins:01], primarily aim to maintain basic network-layer connectivity to a mobile node (MN) as it changes its point of attachment to the fixed network infrastructure. For future wireless Internet access scenarios, an IP–based mobility management framework must not only provide ubiquitous connectivity, but also satisfy the often-divergent QoS requirements of different application sets. In fact, the ability to provision, ensure, and migrate QoS guarantees as an MN moves is an essential component of true multimedia mobility support.

Two principal models, the DiffServ [Blake:98] framework and the Integrated Services (IntServ, a.k.a. resource reservation protocol [RSVP]) [Braden:94, Braden:97] approach, have been proposed for QoS support in the Internet. Of these two, the DiffServ model is judged to be more scalable, as the IntServ model not only requires each intermediate router to maintain flow-specific state, but also needs additional tunneling mechanisms to work across IntServ-incapable intermediate nodes. Both models were, however, primarily designed for the fixed Internet: In DiffServ, the classification and traffic-conditioning policies at the network edge are typically set up offline, while in IntServ, reservation guarantees established via RSVP apply to a specific path and need to be re-established if one of the communication endpoints move.

Since the connectivity from the cellular domain to the Internet backbone is essentially static, traditional DiffServ mechanisms can be used to guarantee performance bounds up to the ingress, or beyond the egress, point of the cellular domain. To extend the QoS guarantee up to the mobile wireless device, it is also necessary to develop reservation protocols to over the wireless link (the "last hop")—a number of intelligent wireless link layers have thus been proposed to provide QoS guarantees over the error-prone wireless channel.

Our focus is, however, on the cellular access domain, which lies in between the wireless edge and the Internet backbone. As different mobile users, potentially with widely varying QoS requirements, continually change their point of attachment to the cellular domain, the traffic loads and the necessary QoS bounds needed on each intra-domain route can exhibit significant fluctuation. Accordingly, the conventional static QoS provisioning mechanisms are not appropriate for the cellular access domain; we require a QoS framework where QoS bounds can be signaled and network resources provisioned in a fairly dynamic manner. As such, we focus on adding two specific functions to the conventional DiffServ architecture:

1. *Dynamic QoS Signaling:* The QoS requirements of different MNs are communicated to the appropriate network nodes, especially as the MN changes its point of attachment.
2. *Resource Provisioning:* Since node mobility makes the traffic load and associated resource requirements on any path dynamic, the framework must be capable of dynamically altering the network resources (such as bandwidth, buffer sizes, etc.) allocated to different traffic classes or flows.

Since the cellular domain can easily have a few thousand (or tens of thousands) concurrent mobile users, such a QoS–aware framework must not, however, a) impose the maintenance of per-flow state information at all intra-domain nodes or b) require resource re-provisioning at each individual node movement.

In this paper, we present our DiffServ-based hierarchical QoS architecture associated with the IDMP. IDMP has been proposed in [Misra:00A, Das:00] as a two-level hierarchical mobility management protocol that reduces handoff latency and the signaling overhead of frequently roaming hosts by localizing most mobility management signaling within a wireless access domain. In IDMP, a specialized node called the mobility agent (MA) provides a stable gateway to the Internet, while subnet agents (SAs) located at the wireless access edge provide the MN an additional local care-of address (CoA) to handle intra-domain mobility. The QoS solution is part of the DMA, which assumes the existence of multiple MAs in the access domain to provide robustness and uses load-balancing strategies to dynamically distribute the MN population set among different MAs.

The DMA QoS solution is based on the use of a BB [BB], a centralized node that controls all resource allocation and admission control strategies inside the wireless access domain. Prototypes and simulation studies have demonstrated the feasibility of using such a centralized solution for provisioning resources and performing admission control in moderately sized domains. The DMA QoS solution was first introduced in [Misra:00B] and consists of the following additional features:

- Additional *IDMP signaling* between the MN, the MA, and the SAs that allows the MN to communicate its QoS guarantees during its initial entry into the domain, and obtain seamless QoS support without additional re-negotiation on subsequent intra-domain movement.
- Additional *coordination* between the MA, the BB, and a new node called the mobility server (MS), which allows different MNs to be assigned different MAs, based on their individual QoS requirements and the network's current traffic-load patterns.
- Dynamic *provisioning* of resources for an MN. Whenever an MN changes its point of attachment, the corresponding MA determines if additional resources are necessary on the new traffic path and subsequently requests the BB to re-allocate intra-domain resources as necessary.

We first present the functional architecture for DMA–based QoS support, including the signaling flow among the constituent functional nodes and the enhancements needed to the IDMP specifications. We then describe our current implementation of the DMA architecture and the validation of our design on our experimental testbed. Performance results indicate that our IDMP–based QoS mechanism is able to support differential performance guarantees for MNs in a scalable manner. Additionally, the performance results demonstrate the rationale for our preference for certain MA–assignment strategies.

The rest of the paper is organized as follows. In the next section, we present an overview of basic IDMP and DMA functionality, as well as review alternative proposals for QoS support in the wireless access domain. In section III, we present the functional components of the QoS solution using IDMP and define the messaging interaction between the component nodes. In section IV, we detail our implementation of QoS–enabled IDMP and the BB, while, in section V, we discuss the performance results obtained on our testbed. Finally, section VI concludes the paper.

II. Review of IDMP and Alternative Intra-Domain QoS Strategies

In this review section, we present a short summary of IDMP functionality and discuss related work in the area of integrating QoS support with IP–based mobility management solutions. In understanding the functional details, it is helpful to distinguish between IDMP (which is a protocol) and DMA (which is an architecture for scalable, QoS–aware mobility management that uses IDMP).

A. IDMP and the DMA Architecture

IDMP is one of several approaches and protocols proposed for introducing a mobility management hierarchy. IDMP is part of a class of solutions, such as mobile IP regional registration (MIP–RR) [Eva:01] and hierarchical mobile IPv6 (HMIPv6) [Soliman:01], which develop an agent-based hierarchy by associating an MN with multiple CoAs and which do not require modifications to the current routing protocols. In contrast, proposals such as cellular IP [Campbell:00A] and HAWAII [Ramjee:00] use a more decentralized approach for intra-domain mobility management. In this approach, the MN is associated with a single CoA that has no locational significance. Host-specific routing entries are established at all intermediate intra-domain nodes to correctly forward packets addressed to this CoA to the MN's current location.

IDMP was introduced under the name TeleMIP in [Das:00] as a two-level hierarchical mobility management protocol for IP–based cellular networks. The DMA architecture [Misra:00A] uses IDMP for intra-domain mobility management; additionally, the architecture specifies the use of distributed MAs and the interaction with multiple global binding protocols. In this approach, networks are partitioned into mobility domains, each consisting of multiple IP subnets. To route packets to the MN's current location, the MN is associated with CoA. The global care-of address (GCoA) resolves the MN's location only to its current domain and hence remains unchanged as long as the MN moves within the domain. The local care-of address (LCoA) identifies the MN's precise point of attachment within the domain and accordingly changes with every change in subnet. To support this hierarchy, IDMP specifies a new functional entity called mobility agent (MA), which is similar to the foreign agent (FA) in conventional mobile IP but resides at a higher level in the network hierarchy. Global updates (to external servers such as mobile IP's home agent (HA) or to correspondent nodes (CNs) are generated only during *inter-domain* mobility and provide the MN's GcoA. In contrast, the MN informs its MA of its new LCoA via intra-domain binding updates whenever it changes its point of attachment. Packets from external nodes are routed (either directly or via tunneling) to the MN's GCoA. The MA intercepts these packets and then tunnels them to the MN's current LCoA. Besides limiting the scope of most binding updates to the local domain, IDMP also provides configurable support for additional mobility-related features, such as fast handoffs and paging (see [Misra:00C]).

B. Related Work on QoS and Mobility Management

Relatively little work has been performed on investigating alternative approaches for combining QoS guarantees with IP mobility support. The problem of using RSVP for mobile nodes was discussed in [Raj:96], which showed how MIP tunneling prevented RSVP RESV and PATH messages from correctly allocating resources along the path from the HA to the MN. Since all packets were encapsulated and tunneled by the HA to the MN, intermediate nodes did not directly see RSVP messages and could not set up appropriate reservations. Moreover, IP–in–IP encapsulation also obscures the inner header; even if the intermediate routers between the HA and the MN did have reservation states set up, they would be unable to classify packets without looking inside the encapsulated payload. Accordingly, it would be difficult to ensure specific QoS guarantees to *a subset of flows* directed toward a specific MN. To solve this problem, Foo and Chua [Foo:00] proposed a modified version of tunneling support in RSVP. In this approach, RSVP PATH and RESV messages are generated by the tunneling endpoints (such as the HA and the FA) to allow intermediate routers to process such messages. In addition, the proposal defines a user datagram protocol (UDP) encapsulation technique that permits the HA to put additional QoS–specific information in the outer header to enable differential treatment by intermediate routers. However, the absence of a hierarchy in MIP for decomposing the resource reservations leads to problems with this approach. Whenever an MN changes its point of attachment (and hence its CoA), new RSVP messages will need to be exchanged along the entire path from the CN to the MN, even though the change may have been confined to the last hop. MRSVP [Talukdar:00] is another approach toward modifying standard RSVP for mobile hosts. In this approach, the MN is assumed to know the set of potential subnets that it may visit during the lifetime of the session; specialized agents on such subnets set up passive resource reservations in advance to account for possible changes in the MN's point of attachment.

A multiprotocol label switching (MPLS)–based solution for QoS in MIP was presented in [Choi:01]. This approach essentially requires the HA to set up a label-switched path (LSP) for MPLS–based forwarding of packets to the MN. In typical implementations, packets addressed to a registered MN are intercepted by the HA and tunneled to the MN's CoA along the fast path through the use of virtual interfaces. The MPLS–based solution however requires additional per-packet classification and labeling functionality at the HA, thereby disallowing use of the fast path and leading to potential performance concerns.

A DiffServ–based architecture for QoS support in mobility domains has been proposed in [Chen:00]. The architecture uses a centralized QoS global server, or QGS, (functionally equivalent to the BB) to manage resources and perform admission control for all MNs in a domain. The QGS also interacts with the ingress nodes, called QoS local node (QLN) located at the ingress edge of the wireless interface to dynamically set up the conditioning and marking functions for individual MNs. However, unlike IDMP's QoS framework, this approach does not separate intra-domain QoS from inter-domain QoS management. Moreover, the solution potentially requires all ingress routers in the domain to parse the inner header in encapsulated packets.

III. The QoS Architecture in DMA

Under the DMA approach, we split the end-to-end QoS management into two distinct parts. The global QoS framework uses the conventional Diffserv/MPLS framework to manage resources between the MA and the global Internet, while the intra-domain QoS framework localizes the dynamic component of resource provisioning within the domain and provides QoS support between the MN and the MA. The architecture requires all inbound (to the MN) and outbound (from the MN) packets to be routed via the MA. Since the MAs are static nodes and handle traffic for an entire population of MNs, global (inter-domain) QoS provisioning can be handled by static resource provisioning and the specification of conventional service-level agreements (SLAs). Intra-domain resource provisioning, on the other hand, is more dynamic (on account of node mobility) and requires enhancements to the base IDMP specifications. In accordance with the DiffServ architecture, the wireless access domain conditions the traffic for the MN at the following two points:

1. For inbound traffic, the MA shapes/polices traffic (before it is tunneled to the current LCoA) to the specified inbound rate.
2. For outbound traffic, the SA shapes/polices packets (before they are tunneled to the MA) to ensure conformance to the outbound traffic descriptor.

A. Functional Layout of QoS-Related Elements

The functional architecture for QoS support is shown in *Figure 1*. The DMA specification includes two new entities, the mobility server (MS), which implements the load balancing algorithm that dynamically assigns MA(s) to an MN, and the bandwidth broker (BB), which dynamically modifies network resources reserved for different traffic classes at each intermediate domain node. In simple terms, the QoS functionality involves the following fundamental steps:

1. The MN must first specify its QoS requirements during its initial registration with an SA in the domain.
2. These requirements are relayed by the SA to the MS, which uses a load-balancing algorithm and other policies to dynamically assign the MN to a specific MA. When the MN registers with the MA (to obtain a valid GCoA), the MA may subsequently modify the QoS assurances.
3. On every subsequent movement within the domain, the MA is responsible for communicating the MN's outbound traffic descriptor and classification policies to the new SA—this obviates the need for additional QoS signaling from the MN at every move.
4. Since the MN's movement changes the path between the MA and the MN, the current resources allocated for the MN's traffic classes along the new path may not be adequate to satisfy the assured QoS levels. In that case, the MA will request the BB to provision additional resources on this new path.

We now provide a brief description of the QoS–related functionality of each of the conventional IDMP elements, as well as the function of the new IDMP elements.

MN: The MN is responsible for indicating its QoS requirements (possibly for different classes) during its initial registration in the domain through extensions specified in the IDMP subnet-specific and intra-domain location update messages.

SA: The SAs essentially behave as edge nodes for the cellular domain. For outbound traffic, the SA classifies packets into different classes (marking) and conditions (policing/shaping) them to ensure conformance to the negotiated traffic descriptor. Since inbound traffic is policed for conformance at the MA, the SA simply forwards such traffic to the MN. The SAs are informed about an MN's traffic profile by the corresponding MA(s) using IDMP control messages. When an MN first moves into a domain, the serving SA also queries the MS to obtain the identity of the appropriate MA(s).

FIGURE 1

QoS–Related Functional Elements of the DMA Architecture

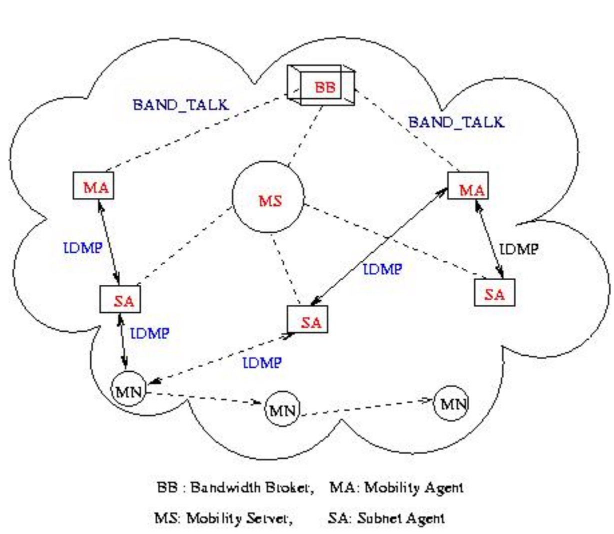

MA: The MA delineates the boundary between the global and intra-domain QoS portions. For outbound traffic, the MA is responsible for conditioning the aggregate traffic to ensure conformance to the global (inter-domain) traffic descriptor. For inbound traffic, the MA must condition packets and mark packets on a per–MN (and per-class) basis to ensure conformance to the specified downstream traffic profile. To perform admission control for a new MN, or to provision additional resources as existing MNs move within the domain, the MA interacts with the BB. *We assume that the MA typically contacts the BB to request bandwidth based on aggregate traffic, rather than issuing reservation requests for individual MNs*. Whenever an MN changes its point of attachment, the MA will need to inform the corresponding SA of the MN's current traffic descriptor and the necessary conditioning actions. In the vast majority of such cases, the MA will not need to interact with the BB, as the aggregate resources will usually prove to be adequate. The MA will interact with the BB only when the current aggregate reservation levels on the new path are not adequate to handle the new traffic load.

BB: The BB is the central entity for admission control and QoS provisioning and maintains information about the resource allocation and consumption at all intra-domain nodes. When the MAs request the BB for additional bandwidth (and other resources) on specific paths to accommodate dynamically varying traffic loads, the BB issues requests to internal domain nodes (including the SAs and the intermediate routers) to dynamically alter the resources allocated to different classes. The BB can also negotiate with BBs in other external domains for provisioning global QoS bounds.

MS: The MS acts as the intelligent control entity that allocates MA(s) to a new MN in the domain. When an MN moves into a domain, the serving SA requests the MS to assign one (or more) MA(s) to the MN. Unlike the BB, the MS is not involved in configuring resources at intermediate nodes but simply in balancing the specified traffic load according to configured policies.

B. Typical Messaging Flow

Figure 2 provides the messaging flow as an MN first enters a new mobility domain in the QoS–aware DMA architecture. The added functionality required the definition of additional IDMP messages and message extensions as follows:

- The *Subnet-RegistrationRequest* message has been extended to include QoS–specific extensions, such as the peak rate and desired packet delay.
- Extensions to the *Intra-domain LocationUpdate* message allow an MN to specify the requested QoS parameters to the candidate MA.
- Extensions to the *Intra-domain LocationReply* message allow an MA to specify the assigned QoS parameters to an MN. This is necessary as the MA may modify the QoS parameters requested by an MN.
- The *UserUpdate* message is sent by the MA, upon an intra-domain movement by the MN, to the new SA to

provide it with the MN's traffic descriptor and conditioning policy.

- A corresponding *UserAcknowledge* message has also been defined for communication from the SA to the MA.

Additional non–IDMP message exchanges are necessary for communication between the IDMP–specific nodes, such as the MA and SA, and the QoS–specific entities, such as the MS and the BB. We shall describe our current implementation of the *MA/SAfl‡BB* communication in section IV.

C. MA Assignment Strategies

As the entity responsible for assigning an MA to an MN registering in the domain, the MS is a key component of the DMA architecture. Different MA assignment algorithms may be appropriate for different QoS differentiation scenarios. In the simplest (and most common) form of QoS support, different users are guaranteed different traffic rates (from a small subset of rate profiles), with uniform treatment for all traffic belonging to a single user. For this model, [Misra:00B] recommends that an MA manage the mobility and QoS requirements of users belonging only to a single traffic class; different sets of MAs are used to handle traffic belonging to different traffic classes. Since an individual MA supports only a single class of traffic, it can apply a uniform header-marking policy for all inbound traffic. Removing the need for differential treatment of different user packets reduces the potential for performance bottlenecks, since experimental evaluation of our implementation (discussed later) shows that QoS–related functionality can significantly add to the packet forwarding latency at the MA. In the more advanced differentiation scenario, different flows belonging to a single user are classified into separate traffic classes; a single user thus subscribes to multiple QoS classes. Even in such a scenario, analysis shows that it may be better for a specific MA to service traffic belonging to only one service class: Not only does this alleviate processing-related concerns, but it also allows an MN to request different mobility support features for different traffic types. Of course, this approach requires enhancements to IDMP, since an MN must now be concurrently associated with multiple MAs (and hence multiple GCoAs). For a detailed discussion of the pros and cons of different MA–assignment strategies and their effect on IDMP messaging flows, the reader is referred to [Misra:00B].

IV. Prototype Implementation

In this section, we describe our prototype implementation of the DMA functional elements. In our current version, we do not have the MS functionality; thus, all requests for traffic belonging to a specific QoS class are routed to a single MA, based on manual configuration of policies at the SA. Our implementation of the MN is obtained by modifying the Linux-based MosquitoNet [MosquitoNet] implementation of MIP; the MA, SA and BB code are based on the FreeBSD platform.

To allow communication between the BB and the MA/SA, we implemented a simple UDP–based BAND_TALK protocol. This is a simple message-exchange protocol where communication takes place through four different message types: request, query, report, and configure. The sequence of messages exchanged by this protocol can be briefly described as follows:

1. When an MA needs to allocate bandwidth on a certain path, it issues an appropriate *request* to the BB.
2. The BB then issues *query* messages to the appropriate intermediate domain routers, asking them for current resource reservation and utilization figures on various links.
3. All the intra-domain nodes then send their status information through appropriate *report* messages.
4. Finally, the BB issues new *configure* messages to the various intermediate routers, as well as to the MA, indicating the successful completion or rejection of a new reservation.

FIGURE 2

Signaling Flow for QoS in the DMA Architecture

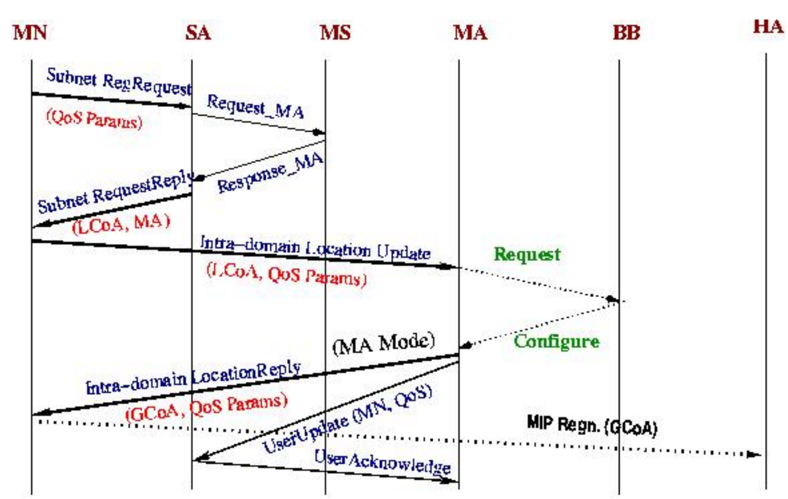

In our current implementation, differential treatment for packets belonging to different classes is implemented via the class-based queuing (CBQ) code (CBQ/ALTQ) from [ALTQ], which implements TOS–based classification and priority scheduling of packets. For our prototype implementation, we implemented just two traffic classes: the default best-effort traffic class and the premium gold traffic class. Service for the gold traffic class is specified through a guaranteed service rate; the service essentially mimics the virtual leased-line service that is developed using the expedited forwarding (EF) per-hop behavior (PHB) [Jacobson:99]. Admission control/provisioning for the gold class is performed by checking that the available bandwidth on the traffic path is greater than the requested rate; traffic conditioners (SA/MA) *shape* the traffic to ensure that the profiled peak rate is not exceeded. We now detail the software architecture of the various DMA functional elements.

A. Mobility Agent

Figure 3 shows the distinct components, and their logical connectivity, in our implementation of IDMP's MA. The main components of the software suite are as follows:

- *MA_Server:* This module establishes the communication between the MA and the BB. It uses a well-known UDP port to communicate with the BB, with retransmissions for reliable data transfer. The MA_Server sends request messages to the BB to provision a new MN (or an MN that has changed its point of attachment). In response to the corresponding configure message from the BB, the module interfaces with the ALTQ/CBQ daemon to update the bandwidth reserved for different TOS classes. To allow the BB to obtain resource utilization from the various intermediate routers (including the MA and SA), the MA_Server daemon must also respond to BAND_TALK query messages that are multicast to a domain-scoped multicast address. Accordingly, it includes a Multicast-Server module that manages the multicast socket interface that responds to multicast messages from the BB.
- *BAND_COLL Agent:* This process calculates the current traffic loads associated with each Layer-2 interface and returns the current bandwidth utilization (for each class) to the MA_SERVER process. The monitoring of traffic loads on each interface is achieved through the netstat system call and a Java-based parser to collect statistics by appropriately parsing the netstat generated data.

In addition to the above components, the MA kernel code has also been upgraded to provide traffic-conditioning functionality. The *conditioner* filters/polices/shapes incoming packets on a per-user basis and is embedded into the kernel, with the use of virtual interfaces to redirect packets to the appropriate conditioning code. New iotcl calls are defined to allow the MA_Server to communicate the traffic descriptor for an individual user (IP address) to the conditioner and to allow the MA_Server module to update the CBQ/ALTQ daemon with modifications in the bandwidth reservations.

B. Subnet Agent and Intra-Domain Routers

The SA software component is very similar to that of the MA; both the SA and the MA are essentially edge nodes in the DiffServ–enabled cellular access domain and thus perform similar actions such as traffic classification and conditioning. The SA component thus not only contains the ATLQ/CBQ code, but also the functionality of the BAND_COLL agent and the conditioning agent, as well as the kernel-layer conditioner code. Furthermore, as shown in *Figure 2*, the SA and the MA are responsible for exchanging IDMP messages that inform the SA of the traffic descriptor corresponding to a specific MN.

For our prototype implementation, each intermediate router in the cellular domain was configured to run the SA_Server code to allow such nodes to respond to query messages from the BB and also to receive configure messages that

FIGURE 3

Organization of MA's Software Modules

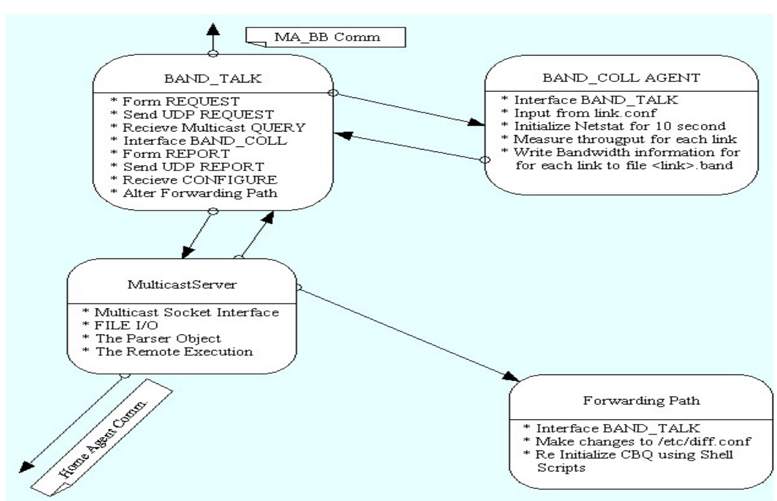

adjust the resources reserved for different service classes.

C. Bandwidth Broker

Fundamentally speaking, the BB has several different functions to fulfill in the QoS assurance process:

- It must receive new configuration requests from the MA(s) and then respond to them.
- The BB must query all intra-domain nodes to determine current resource reservation and traffic utilization levels for different traffic classes.
- Given the set of concurrent resource requests in the domain, the MA must be able to determine the appropriate resource provisioning required at all intra-domain nodes. If adequate resource provisioning is not possible, the BB must reject such reservation requests.
- Finally, the BB must be able to communicate any changes required to the current resource reservations for various traffic classes back to the intermediate routers in the domain.

For portability, the BB functionality was implemented as a Java-based application. The software architecture of the BB consists of the following components:

- The *BAND_TALK* client: This process interacts with the MA to receive and respond to new reservation requests, as well as with the SA and other intermediate routers to obtain utilization information and send configuration requests. As stated earlier, the BB collects usage statistics by sending BAND_TALK query messages to a multicast address and provides new configuration parameters to domain nodes via configure messages. On receiving report messages from various domain nodes, the BB spawns separate threads that store and process the reported link utilization information. To store this information efficiently, the BB uses a synchronized array object that allows consistent access to data across multiple application threads.
- The *BAND_FLOW* process: This process is at the heart of the BB functionality, since it implements the admission control and resource-provisioning algorithm that must be invoked for every new BAND_TALK request from an MA. The admission control for a best-effort user is simple: By definition, best-effort flows are always admitted. To perform admission control for a gold class (priority) user, we have implemented a simple max-flow–based algorithm described next in *Figure 4*. The algorithm assumes that the entire cellular domain is represented as a graph, with vertices corresponding to the intra-domain nodes and edges corresponding to the traffic links between such nodes. Link i is assumed to have a maximum capacity of C_i.

The algorithm is very simple to understand. Whenever a new request arrives at the BB, the BB first obtains the least-loaded path from ingress to egress, i.e., the path that maximizes the residual bandwidth on the bottleneck link. Subsequently, it computes whether the requested resources can be reserved on that particular path. The first path of this determination is straightforward: If the reserved rate for gold class traffic can be increased by the requested rate without violating the capacity constraint, then it is always accepted. As a secondary heuristic to improve network utilization, we also admit the new traffic request if such provisioning would not cause the current utilization levels to exceed the capacity by more than *15 percent*. (This is based on our observation that multimedia traffic often does not utilize the reserved peak rates; hence, some statistical gain is possible.) We do not make any optimality claims about this heuristic mechanism; indeed, our focus here is not on the exact admission control scheme but on the associated signaling architecture. Thus, an operator that wishes to perform worst-case admission control for the gold class would clearly run the Ford-Fulkerson procedure with link weights set to $Resid_i = C_i - Resv_i$ and accept the request only if $Req + Resv_l < C_l$ alone.

V. Experimental Set-Up and Implementation

To study the performance features of our implementation, we set up a testbed as shown in *Figure 5*. Two independent IEEE 802.11-based subnets were defined, and the SAs (not shown in the figure) were co-located with the wireless local-area network (LAN) access points. As the MN moved across

FIGURE 4

Admission Control Algorithm for Gold Class

- ❖ Assume that a new request is made for *Req* capacity from node *A* (ingress) to node *B* (egress).
- ❖ Use the BAND_TALK protocol to obtain the current utilization, $Util_i$, as well as the resource reservation for the gold class, $Resv_i$, for each link i in the domain.
- ❖ Set up a graph ***G*** representing the domain topology with the weight of link i equal to the residual capacity $Resid_i = C_i - Util_i$.
- ❖ Run the Ford-Fulkerson algorithm to determine the max-flow path P in G from A to B.
- ❖ Obtain the bottleneck link $\boldsymbol{l}$ (one with the least residual capacity) in path P using $\boldsymbol{l = \arg\min j: Resid_j ; j\ e\ P}$

If $Req + Resv_l <= C_l$, then accept request;
If $Req + Util l <= 1.15 * C_l$, then accept request;
Else reject request.

- ❖ If request is accepted, then upgrade $Util_i$, $Resv_i$ for all links[1] on path P.

the LANs, it would be assigned a new LCoA, although its MA would not change (thus, the entire network was treated as a single domain). The QoS–related code was implemented on the MAs and the SAs. Depending on the traffic path, an MA might also act as an intermediate router. For the purpose of comparison, we also deployed a publicly available version [CIP:01] of cellular IP on our testbed and performed identical experiments. It should be noted that, in contrast to our IDMP implementation, this implementation of cellular IP lies in the user space. Experiments with various settings of the reserved rate for the gold class, and concurrent gold and best-effort flows, verified that our architecture was able to ensure preferential queuing for premium flows and insulate their performance from the best-effort traffic load.

To further study the effects associated with our QoS implementation, we performed experiments where only a single active traffic flow (UDP or TCP) was present in an otherwise unloaded network. We carried out experiments by first classifying the active flow as gold traffic and enabling the QoS feature of DMA, and subsequently by classifying the active flow as best-effort and disabling DMA's QoS support. In the absence of any cross traffic, the effective throughput and loss characteristics should ideally be independent of whether the flow is classified as gold or best-effort; any differences in performance results can really be attributed to the processing and forwarding overheads associated with our QoS implementation. We now report on experiments that attempt to investigate the following:

1. The impact of QoS processing on the handoff latency in IDMP
2. The effect of the DMA QoS features on the forwarding latency at various intra-domain nodes and consequently the overall throughput for TCP flows.

A. Handoff Latency

To first evaluate the handoff latency under IDMP, we injected a periodic UDP traffic stream from the CN to the MN, varying the intra-packet gap allowed us to adjust the transmission rate. By observing the sequence of packets lost during a manually triggered handoff, it is possible to bound the potential handoff delay. Under the assumption of constant one-way delays, a handoff loss of exactly N packets, each generated at a uniform interval of X ms, corresponds to an upper bound of $(N+1)*X$ ms and a lower bound of $(N-1)*X$ ms on the handoff delay. *Figure 6* plots the number of packets lost during handoff against the inter-packet gap, with both DiffServ–enabled (gold class) and DiffServ–disabled (best-effort class). For comparison purposes, we include the corresponding results observed on our testbed with the publicly available implementation of cellular IP hard handoff [Campbell:00B,CIP:01].

In our topology, we can see that the IDMP handoff is bounded by ~150 ms; additional IDMP fast handoff mechanisms proposed in [Misra:00c] should reduce this transient loss and handoff delay significantly. More importantly, note that the introduction of QoS mechanisms actually decreases the handoff loss somewhat; the processing and forwarding delay of QoS–related processing can actually lower the number of packets incorrectly forwarded to the old point of attachment. The graph also indicates that the *QoS–related signaling in IDMP does not significantly increase the handoff latency.*

We can also see that, even after incurring the overheads associated with QoS–related processing, the IDMP handoff delay is lower than that obtained with cellular IP. While we suspect that this is largely due to the user-space implementation of cellular IP, our results indicate that IDMP performance is, if nothing else, at least as *competitive* as alternative intra-domain mobility management solutions.

B. QoS–Related Processing Overhead

To quantify the forwarding overheads incurred due to the QoS–related processing in DMA, we subsequently observed the throughput of a single TCP flow as a function of the MSS

FIGURE 5

Topological Layout of Testbed

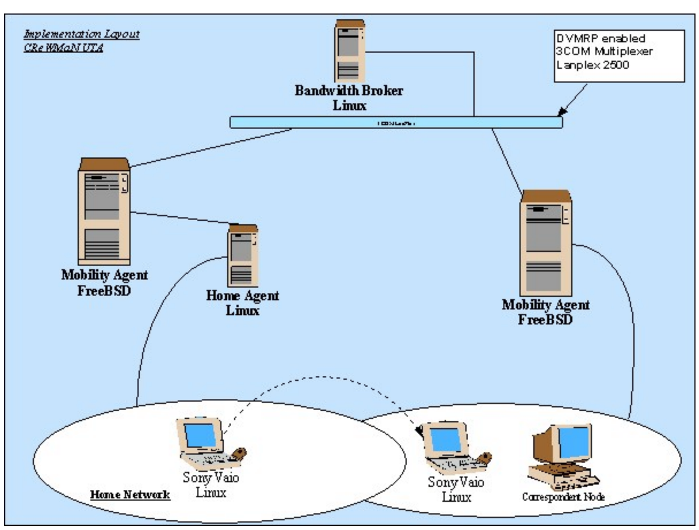

FIGURE 6

Packet Loss during Handoff Transient

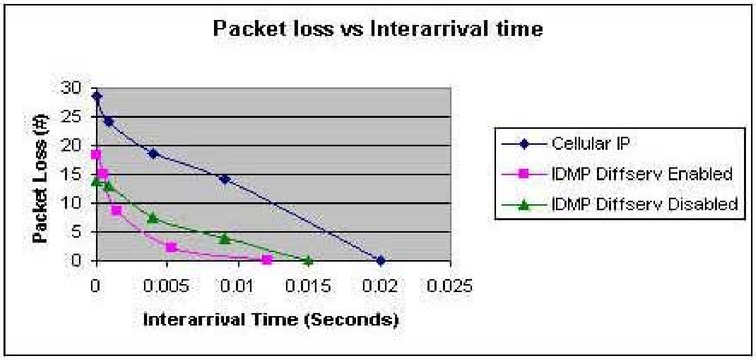

of the TCP connection, first with DiffServ enabled and TCP classified as gold traffic, and then with DiffServ disabled and TCP classified as best-effort traffic. This set of experiments involved no handoffs, as the MN was kept stationary. Clearly, functions such as classification and conditioning (on a per-user basis) at the MA and policing/shaping (on a per-class basis) at all intermediate nodes, including the MA and the SA, contribute to forwarding overhead. Since the overheads typically occur for every individual packet, the difference between the DiffServ–enabled and DiffServ–disabled cases is accentuated if the number of transmitted packets is larger. Accordingly, we would expect throughput differences to be larger for smaller MSSs, since a larger MSS implies that data is transferred in larger chunks (and consequently a lesser number of individual packets). This behavior can indeed be observed in *Figure 7*, which shows the variation in throughput as the MSS is varied between 256 and 2,400 bytes. The figure shows that the Diffserv–processing overhead can indeed be significant, since the throughput drops from ~2.7 Mbps to ~2.4 Mbps when the QoS–related processing is activated. However, this difference diminished as the MSS values correspond to larger (and more realistic) values. The sharp drop in the throughput beyond an MSS of 1,500 bytes occurs because the MSS exceeds the Ethernet link maximum transmission unit (MTU). The behavior is then dominated by the effects of packet fragmentation, which increases the layering overhead and degrades the effective throughput. The graph in *Figure 7* illustrates the potential for bottlenecks at the MA and supports our suggestion for MA–allocation architectures (as described in section III-C) that minimize the classification and marking functionality necessary at the MAs.

C. Handoff and Processing Overhead Trade-Offs

To further investigate the trade-offs between the processing overheads associated with traffic conditioning/policing and the mobility-related signaling in IDMP, we studied the throughput of a TCP connection as a function of the handoff

FIGURE 7

Impact of IDMP's QoS Processing on TCP Throughput

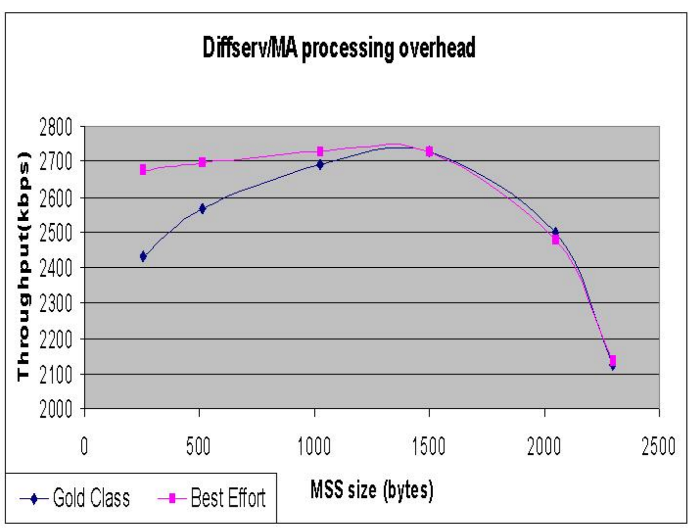

rate. *Figure 8* shows the plots for IDMP and cellular IP as the number of handoffs per minute is varied from *0* to *10*. For IDMP, we experimented with both DiffServ–enabled (gold-class traffic) and DiffServ–disabled (best-effort traffic) scenarios. The plot shows that for small handoff rates, the DiffServ–disabled TCP flow obtains a higher throughput than an equivalent DiffServ–enabled TCP flow. However, as the handoff rate increases beyond six handoffs per minute, the DiffServ–enabled TCP flow obtains a higher throughput. In fact, the performance of DiffServ–enabled TCP is relatively independent of the handoff rate. This can be explained by realizing that, at higher mobility rates, the throughput is affected more by interruptions due to handoff latencies than due to per-packet processing overheads. In fact, the processing delays caused by CBQ and classification functions at the MA serve to reduce the number of misdirected and lost packets during the handoff transient, thereby averting a significant reduction in the TCP throughput. The plot also shows that, even with the forwarded overheads associated with QoS processing, our DiffServ–enabled IDMP implementation is able to achieve better throughput than the cellular IP implementation (which does not suffer from any QoS–related overheads).

VI. Conclusion

In this paper, we presented architecture for supporting QoS guarantees to mobile nodes under the DMA framework. Our proposed solution tackles the signaling and dynamic provisioning challenges in the cellular access domain that result from the potential mobility of hosts *within* the domain. We have shown how the IDMP two-layer intra-domain mobility hierarchy can be augmented to allow the MN to signal its QoS requirements during the initial registration in the domain. The architecture proposes a load-balancing strategy, where each such MN is dynamically assigned to potentially different MAs through an MA–assignment algorithm executed by a centralized MS. We avoid the need for subsequent QoS–related signaling during intra-domain movement by having the MA dispatch the MN's traffic and QoS profiles to the new point of attachment (SA), which can then condition outbound traffic appropriately. To provision resources between the MA and the MN's current point of attachment, we have used the centralized DiffServ–based BB architecture, which is responsible for signaling all intra-domain nodes with the appropriate reservation levels for different traffic classes. We have presented a simple, yet effective, admission control policy for a basic differentiation scenario involving two classes of service—one peak-rate controlled premium and the other best effort.

By experimenting with our prototype implementation on our testbed, we have developed an understanding of the performance and implementation issues. Our experiments show that the handoff latency does not differ appreciably in the presence or absence of QoS control in IDMP, indicating that our signaling framework does not incur significant latency. However, throughput studies show that activating the QoS–specific functions do result in somewhat lower than optimal throughput due to the forwarding latencies associated with QoS processing at the MA. Accordingly, we recommend architectures where the classification and marking functions at the MA are kept to a minimum, thereby alleviating a potential forwarding bottleneck.

Several pieces of our architecture need further development and investigation. For example, the MS still needs to be implemented, especially with suitable algorithms for load balancing across candidate MAs. Moreover, our architecture currently requires *both inbound and outbound traffic* to be relayed via the corresponding MA. In a more optimized framework, outbound traffic need not be relayed through the MA but can be directly routed to the domain egress routers. We need to develop protocols that allow such asymmetric packet forwarding. Security and authentication, especially in the context of QoS support, is another important area. For example, we need to investigate techniques

FIGURE 8

Impact of Handoffs and QoS on TCP Throughput

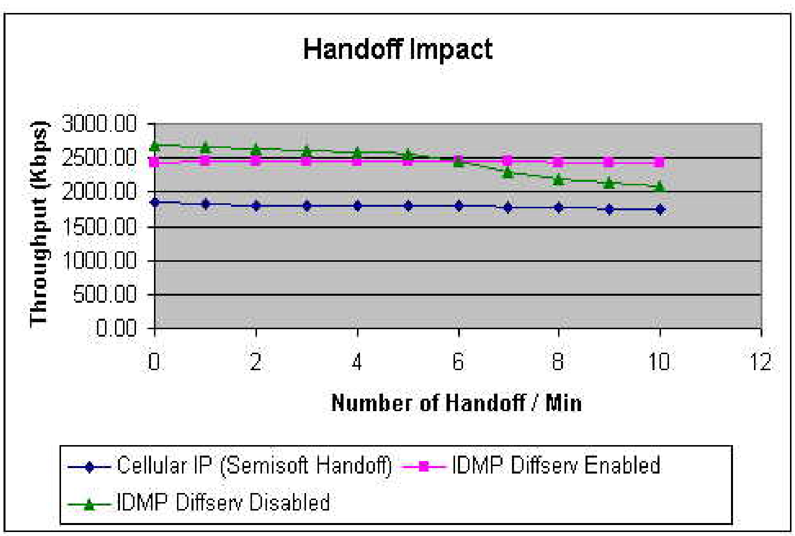

for the secure transmission of UDP–based BAND_TALK messages and their impact on the signaling overhead. Finally, we are also working on evaluating the relative performance of our architecture against implementations of other mobility management solutions, such as NUS mobile IP with RSVP tunneling and HAWAII, on our testbed.

Acknowledgements

The authors would like to thank Muni Bajpai and Wei Wu of the University of Texas at Arlington for their help in setting up the testbed and running the experiments. The work of Sajal Das was supported by an NSF grant, numbered EIA-0115885.

References

1. [Perkins:01] C. Perkins, "IP Mobility Support for IPv4, revised," IETF, Work in Progress (September 2001), draft-ietf-mobileip-rfc2002-bis-08.txt.
2. [Blake:98] S Blake, D Black, et al., "An Architecture for Differentiated Services," RFC 2475 (December 1998).
3. [Braden:94] R. Braden, et al., "Integrated Services in the Internet Architecture: An Overview," RFC 1633 (June 1994).
4. [Braden:97] R. Braden, et al., "Resource ReSerVation Protocol (RSVP): Version 1 Functional Specification," RFC 2205 (September 1997).
5. [Misra:00A] A. Misra et al., "IDMP: An Intra-Domain Mobility Management Protocol Using Mobility Agents," Work in Progress (July 2000), draft-mobileip-misra-idmp-00.txt.
6. [Das:00] S. Das at al., "TeleMIP: Telecommunication Enhanced Mobile IP Architecture for Fast Intra-Domain Mobility," *IEEE Personal Communications Systems* (August 2000).
7. [BB] "Internet-2 Qbone Bandwidth Broker," available at http://www/merit.edu/working/groups/i2-bone-bb/.
8. [Misra:00B] A. Misra et al., "Integrating QoS Support in TeleMIP's Mobility Architecture," Proceedings of IEEE International Conference on Personal Wireless Communications (ICPWC) (December 2000).
9. [Eva:01] E. Gustafsson, A Jonsson, and C. Perkins, "Mobile IP Regional Tunnel Management," Work in Progress (September 2001), draft-ietf-mobileip-reg-tunnel-05.txt.
10. [Soliman:01] H. Soliman et al., "Hierarchical MIPv6 Mobility Management," IETF, Work in Progress (July 2001), draft-soliman-mobileip-hmipv6-04.txt.
11. [Campbell:00A] A. Campbell et al., "Cellular IP," Work in Progress (January 2000), draft-ietf-mobileip-cellularip-00.txt.
12. [Ramjee:00] R. Ramjee et al., "IP micro-mobility support using HAWAII," Work in Progress (July 2000), draft-ramjee-micro-mobility-hawaii-01.txt.
13. [Misra:00C] A. Misra et al., "IDMP–based Fast Handoffs and Paging in IP–based Cellular Networks," Proceedings of IEEE 3GWireless 2001 Conference (May 2001).
14. [Raj:96] B. Rajagopalan, "Mobility and Quality of Service (QoS) in the Internet," MoMUC, 1996.
15. [Foo:00] C. C. Foo and K. C. Chua, "Implementing Resource Reservations for Mobile Hosts on the Internet using RSVP and Mobile IP," Proceedings of IEEE Vehicular Traffic Conference (VTC) (2000).
16. [Talukdar:00] A. Talukdar, B.R. Badrinath, and A. Acharya, "Integrated Services Packet Networks with Mobile Hosts: Architecture and Performance," *Journal of the Wireless Networks*, Volume 5, No.2 (March 1999).
17. [Choi:01] J. K Choi et al., "Extension of LDP for Mobile IP Service through theMPLS Network," IETF, Work in Progress (February 2001), draft-choi-mobileip-ldpext-01.txt.
18. [Chen:00] J. Chen et al., "A QoS Architecture for Future Wireless IP Networks," Proceedings of Parallel and Distributed Computing and Systems (PDCS) Conference 2000 (November 2000).
19. [MosquitoNet] MosquitoNet Project: "Mobile IP Resources," Technical Report (1998), http://gunpowder.Stanford. EDU/mip/ resource.html.
20. [ALTQ] Sony Corporation, "Code for Alternate Queuing altq," http://www.csl.sony.co.jp/person/kjc/kjc/software.html.
21. [Jacobson:99] V. Jacobson, K. Nichols, and K. Poduri, "An Expedited Forwarding PHB," IETF, RFC 2598 (June 1999).
22. [CIP:01] Columbia University Comet Group, "Cellular IP v.1.1 for Linux" (September 2000), http://www.comet.columbia.edu/cellularip.
23. [Campbell:00B] A. Campbell et al., "Design, Implementation and Evaluation of Cellular IP," *IEEE Personal Communications Systems* (August 2000).

Notes

1. In addition to implementing such an algorithm, the network must also have the technological capability to ensure that packets from A to B actually follow the designated path P. This can indeed be achieved by technologies such as MPLS, where explicit paths can be set up, and is orthogonal to our QoS discussion.

Streamlining Operations for Service Providers in the Next-Generation Network Infrastructure

Olivier Daures
Product Manager
Compaq Computer Corporation

Introduction

This paper addresses operations from the point of view of a service provider rather than a vendor, concentrating on the new trends in next-generation networks, in particular wireless mobility and the growing role of IP networks, which bring new challenges to the service providers in terms of operations support systems (OSS).

There are three sections to this paper. The first section discusses services; the second, challenges for service providers in the years to come; the third, OSS requirements.

Services in the Coming Years

Real convergence over Internet protocol (IP) is extending IP everywhere, and everything nowadays tends to support IP. It is predicted that within the next few years, even the radio part of Universal Mobile Telephone System (UMTS), a wireless mobile standard, will be supported by IP. This is not the case now, as the radio part is supported by asynchronous transfer mode (ATM). Another example is low-level optical layers, which tend to support core router connectivity directly. This move forward will continue.

There are clearly two kinds of services worth distinguishing: Internet access (and everything around that) and virtual private networks (VPN) (a service that will increase dramatically in the years to come). It is very important for today's OSSs to be able to support both public and private services. The Internet itself is a large public network, but private IP–based networks, such as those that are made for third-generation (3G) networks and for voice over IP (VoIP), promise to grow. Convergence will affect services because the same services will be provided through a mesh of various access, network, and content providers. Mobility over 3G networks means wireless operators will both compete and cooperate with Internet service providers (ISPs) and network service providers (NSPs). The telecommunications industry is using mostly fixed lines at this point. But in the years to come, a new generation of devices and protocols will bring about value-added wireless mobile services, which will be made available through tiny hypertext markup language (HTML) or new versions of the wireless application protocol (WAP).

VPN Services

VPN services will be based mostly on differentiated services (DiffServ) and multiprotocol label switching (MPLS), providing new features such as closed user group (CUG) and class of service (CoS). IP VPNs are taking over from overlay VPNs, where point-to-point connections, such as frame-relay ATM, support intranets and extranets for business-to-business (B2B) support. Other services on top of VPN include VoIP and IP security (IPSec).

Also, with VPN, products, such as bandwidth on demand (based on resource reservation protocol [RSVP]), will allow for video streams and multicast broadcasting. This, in turn, means things like medical images can be shared in real-time between hospitals. VPNs will enable various value-added services (VAS), such as storage-area networks (SAN) and remote content servers for Web hosting, as well. Furthermore, wireless access will be an important feature in general packet radio service (GPRS) and 3G networks. No longer will laptop users have to make an extra connection. When business users have the opportunity for wireless mobile access, they will buy it; and what is interesting is that they will buy it from wireless operators, not from the current NSPs or ISPs. That may change, however. The ISPs will have to provide that access, so they will have to deal with the wireless operations to find agreements to provide such access.

3G Services

Third-generation services include high-speed data mobile access services in which the 3G network is used purely as access—it can be an access in itself to the Internet, for example. But 3G networks will bring specific services, such as electronic payment (e-payment) and location services, as well. Those application services will mostly be enabled by IP. Other 3G services will rely on voice and value-added voice services.

Architecture for 2.5/3G Services

The classical architecture for the 2.5G services is shown in *Figure 1*. Such 2.5G services include GPRS and UMTS, connecting mobile terminals to the radio network and onto the core backbone, where IP transfers the traffic to the services network. From there, users can be connected to networks of other service providers.

Examples of Mobility

Examples of mobility and its unique offerings abound. *Figure 2* shows how users, with certain handsets, can drag and drop for WAP, accessing HTML to create a page on their telephone set. This capability should prove to be of great interest to most consumers owing a WAP–enabled mobile phone. Another example of specialized mobile applications, as illustrated in *Figure 3*, is downloads with personal digital assistants (PDA). For instance, if an end user is walking down the street and passes a movie theater, that person might want to get some information about one of the movies playing. With a PDA and appropriate mobile access, the user will be able to effectively gather this kind of information. Indeed, it becomes staggering to think about all the potential services that can be and are being made available to consumers today.

Managing More Complexity

Given the shifting landscape in the telecommunications industry, service providers today must be able to manage change and increased complexity effectively (see *Figure 4*). As evidenced by the break-up of AT&T, companies can and must be split up in order to successfully specialize certain domain content providers. Access network operators and core transport network operators today have created a very complex scheme that is a challenge to business support systems (BSS) and OSSs.

Challenges for Service Providers

Service Assurance

Service providers must, first and foremost, reduce the cost of managing their network. They must increase service value by facilitating service level agreements (SLA) or service objectives (SO) for specific services and customers, especially for VPNs. Once SLAs are in place, providers have to define the parameters of their service guarantees and monitor the level of service being offered; ideally, these providers might even be able to exceed customer expectations by proactively resolving problems. Service providers must optimize the use and repair of the network for most revenue-generation services to strengthen service assurance capacity.

Service Delivery

Equally important as establishing this service assurance, service providers must also concentrate on reducing the lead-time to connect. That is, they must shorten the time necessary to deliver service to a customer. Automated service activation is crucial for this.

Time to Market

In addition, a service provider must work to reduce time to market so that it can be among the first to deliver a given new service. Doing this, however, can place an enormous amount of strain on the BSS–OSS chain. For example, a provider might have to put VoIP service on VPNs in five months' time. But that provider must realize the tremendous, if slow and progressive, pressure and strain that such an endeavor will be causing the BSS–OSS chain.

OSS Solution Requirements

OSS solution requirements include service assurance and service fulfillment. Service assurance includes such features

FIGURE 1

2.5G/3G Services

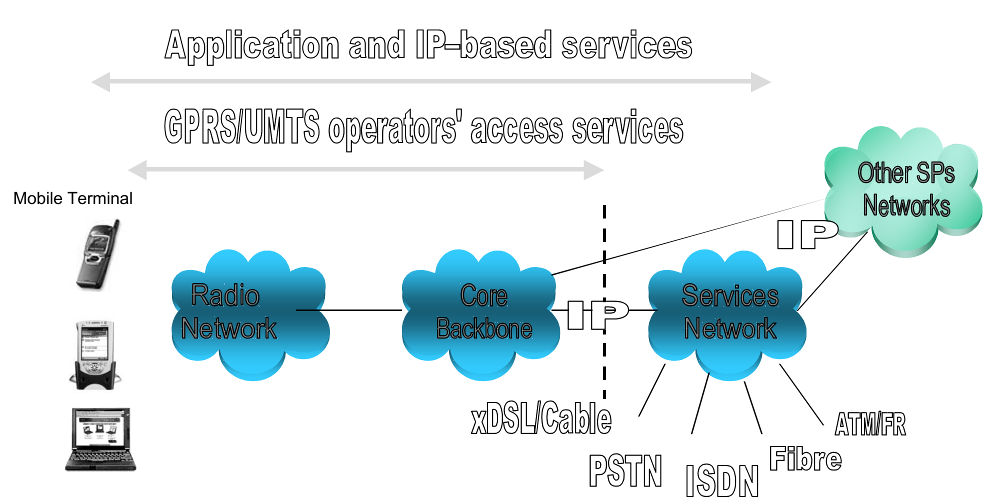

FIGURE 2

Drag and Drop for WAP

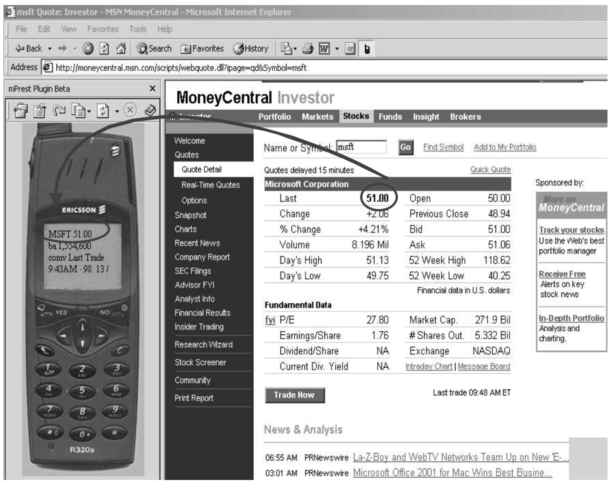

FIGURE 3

Downloads

fault and trouble detection, performance guarantees, service impact analysis, service-level monitoring, and SO monitoring. Taken together, such features might be considered a commodity. However, this is not the case right now because convergence has constraints to consider. In terms of service fulfillment, which entails more than mere service activation, service providers must deal with many challenges, such as the whole provisioning process and workforce management. Other service-fulfillment issues include configuration and inventory management and subscriber and end-user management (something that is forgotten from time to time). New on the horizon is mobile-device management with wireless services.

In terms of architecture, OSS solution requirements must account for cross-domains, such as IP, radio, and other access networks. OSS solutions must accommodate multiple vendors because nowadays it is unlikely that service providers will take everything from one vendor. A multivendor environment means having to manage disparate elementsæa requirement that also means any solution must be scalable. If not, it will be obsolete in a few years. In the IP world, many solutions came from the enterprise market. It is clear that availability and scalability of the system is key for many service providers; and it is very challenging for some OSS vendors to adapt to this kind of environment. It is very important, in fact, for service providers to take into account from the very beginning a solution that is cross-domain, multivendor, scalable, and highly available.

Key Service-Assurance OSS Requirements

In terms of key service-assurance OSS requirements, it is important for service providers to be proactive rather than reactive. To do this, they must monitor services, SOs, and SLAs before a problem occurs. Service providers need a complete end-to-end view, from the end user to the applications. For example, for UMTS, service providers must monitor the radio access part (air interface), the switched network, the core backbone, and the services network(s) of the applications. When a problem arises, service providers must be able to know where the source of the problem is, and they must be able to understand what they have to do to fix it. Service is not just an application on a server somewhere. It goes through the network, or several networks; and when a problem occurs at the end user side, it is very important to know exactly where it is.

OSS service assurance has to be customer oriented. Service providers need to have information in their OSS not only about their services but also about their customers. Service assurance, now more than ever, also needs to support the diverse business relationships that exist in the world today.

Key Service-Fulfillment OSS Requirements

In terms of service-fulfillment requirements, OSS needs to handle convergence for services and technologies. A good example is GPRS, which is perhaps more widely known in Europe than in the United States. Basically, GPRS provides radio access to deliver IP connectivity over the Global System for Mobile Communications (GSM). In GSM, users are managed in a home location register (HLR), which is a completely separate system from the IP subscriber-management system. Does the user have to be managed via the GSM or the IP solution? In dealing with this convergence, issues such as this one are indicative of the types of questions that providers will have to deal with on a daily basis.

Another aspect of service fulfillment is hierarchical support of services and components. Whatever service providers release in their service, they never know if the service will or will not be used as the basis for another service. When service providers delivered frame relay access a decade ago, they did not imagine how that could be used for the IP networks. So OSS requirements demand flexibility to build things on top of the others.

Auto-provisioning is also a key area, one that comes mainly from the wireless world. It is very important for users to be

FIGURE 4

Managing Complexity

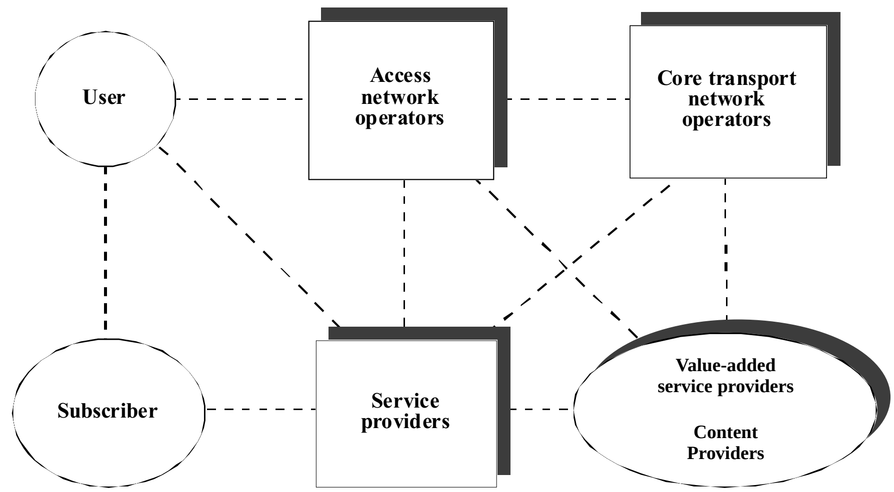

*Source: UMTS Forum

able to change services themselves, if possible with their handsets in hand. Users will eventually demand such a service, and service providers will have to comply with SLA and quality of service (QoS) support.

OSS service fulfillment must be workflow driven and take into account service provisioning of the various service elements and workforce management. For example, in VPN service fulfillment, key actions are delivered by human beings, not by systems. Service providers must follow the whole process.

The architecture should be repository based, mainly because there are too many components to deal with using any other approach. Enterprise application integration (EAI) and flexible formats, such as extensible markup language (XML), are important as well. And it is usually a good idea to have a steady plug-in capacity for applications. What is often challenging is trying to implement a "Big Bang" approach: the Big Bang approach generally does not work because the legacy environment is too complex for a large service provider; plug-in capacity allows for a steady way of dealing with this, though. Finally, Web access is another (perhaps obvious) requirement as well.

When Power Management Becomes a Power Struggle

Dennis DelCampo
Vice President, Customer Fulfillment, North America
Marconi Communications

Power Is the Starting (and Stopping) Point

Dependable power is the starting point for every network system and device, and continuous connectivity is vital to the profitability of any network—be it voice, video or data, copper, fiber, or wireless. Research studies[1] indicate that more than 45 percent of accidental data losses are attributable to power anomalies. Other studies[2] disclose that 50 percent of organizations account for downtime losses from power outage at $1,000 per hour, while 9 percent put their loss costs at $50,000 per hour or more. That is why many Internet-based companies, financial institutions, and other service providers are demanding higher-power system reliability, as high as 99.9999 percent ("six nines").

Growing and Changing Business

At the same time that network operators are striving for greater reliability, the rapid growth and restructuring of the telecom industry has placed additional demands on the distribution of hard-to-acquire capital and the utilization of technical personnel.

This has led to a major philosophical shift—from ownership and maintenance of networks by carriers and service providers to renting from network operators. This new direction is allowing the carriers and service providers to focus their resources and attention on providing an ever-expanding menu of customer services.

Paralleling this shift is a loss of network maintenance expertise through retirement and the discontinuation of the extensive training programs that Western Electric and the Bell companies once conducted—while large amounts of legacy equipment are still functioning in central offices (COs) and other installations that may have been in place for 50 or more years.

The Solution

Obviously, the network won't go away; indeed, it will continue to grow. And new, smarter ways are evolving to maintain it. The transition from reactive maintenance to preventive maintenance to predictive maintenance is a clear indication that this change is under way. Today, telecom industry leaders no longer look upon predictive maintenance as an expense; it is considered a *strategic investment in availability assurance*.

At the forefront in satisfying the demand for increased network dependability are dedicated network service companies with the skills, technology, and global presence to monitor, maintain, and manage telecom power from network operation centers (NOCs) and tech service installations strategically located around the world.

When the network owner cannot or does not choose to handle all or some part of the network-management tasks in-house, the owner can outsource specific tasks or the total network power management activity to a network service company.

Technology in Place

To monitor power systems at thousands of manned and unmanned sites without actually being there, network service companies rely on computerized data-collection equipment and telecommunication links to their NOCs. This data-collection equipment can gather data (including proprietary legacy communications) from a virtually unlimited number of shunts, probes, and sensors installed at the site to monitor equipment performance and site environment conditions. Basically, any condition that can be transformed into a voltage or current can be monitored.

Connectivity to the NOC can range from dial-up modems to switch provisioning to the Internet in order to carry data to a different area and transmit it to the NOC. The reams of data generated and transmitted to the NOC are handled with a variety of software, including event managers, element managers, performance-management packages, tracking packages, and database engines; and the data are transmitted using simple network management protocol (SNMP), a protocol that provides the framework for managing power devices.

How Service Company NOCs Work

The service company's NOCs (see *Figure 1*) are normally staffed by analysts who are very skilled, highly trained individuals that are able to operate in a multivendor environment (can deal with all brands and kinds of power equipment and ancillary equipment as well). The service company itself will normally have a historical understanding of power-equipment performance. Often, the service company may be affiliated with a manufacturer of power equipment.

Through the use of its remote network-management platform, the NOC's element-manager software talks to all attached devices, "polling" the data-gathering equipment installed at the remote site with a regular frequency determined by the network owner. The network owner also works with the service company to establish specific readings or equipment status levels as alarm thresholds.

Major, Minor, and Predictive Alarms

Alarms may be designated as major alarms requiring emergency attention (immediate dispatch of service personnel to the site) or minor alarms that may be deferred for resolution until the next business day. The network owner may also choose to establish "predictive" alarms for further evaluation. Predictive alarms indicate that the likelihood of equipment failure is imminent and that rebuild or replacement should be considered before service interruptions occur. Monitoring and tracking changes in battery string voltage, float current, and conductance, for instance, will enable end-of-life predictions to be made.

When a fault condition is detected, a NOC analyst will diagnose the situation and isolate the trouble source according to a "rules-based" workflow process. The analyst may also confer with technical experts on the system or equipment involved. When the fault is isolated, the NOC analyst will notify the network owner's designated contact if the service company is only responsible for fault *detection*.

If the service arrangement includes fault *reaction* as well, the NOC analyst will attempt to clear the fault remotely if possible. If the condition requires on-site repair and restoration, the analyst will notify the network owner's designated repair contacts within a predetermined time frame. Or, if the service company is also responsible for on-site repairs, the NOC analyst will dispatch the appropriate service personnel from the nearest service center or contact a third-party service provider if special expertise is required.

In addition to power system fault detection and reaction, the NOC, if desired, will monitor any other designated devices at the site, including gen sets, A/C, batteries, room alarms, and other ancillary equipment. Events and data logged by the NOC can then be used for performance reporting and network-planning activities, including energy procurement, power-system coordination, power-capacity planning, network-reliability improvement, and other critical network-management functions.

The Choice Is Yours

Whether you choose to outsource your complete network service needs on a round-the-clock basis, outsource specific tasks, or rely on the service provider for after-hours and weekend support, the use of a network service provider (NSP) will augment your in-house service capabilities or replace lost expertise.

In any case, it will help you to dedicate your personnel and efforts to the core business of providing revenue-producing services and managing your customers, instead of managing your network.

Notes

1. Infonetics Research, Inc.
2. Yankee Group

FIGURE 1

Marconi NOC, Warrendale, PA

Securing Broadband Networks

Sandeep Dhameja

Senior Network Security Consultant

Divine, Inc.

The Business Challenge

The challenge of securing networks is rapidly gaining momentum with the increase in e-commerce. Companies are quickly realizing that to stay both competitive and successful, e-business *has become* their business and that the resources to enable their business, namely the computers and the networks, are information assets. GartnerGroup says that business-to-business (B2B) e-commerce transactions will generate $7.3 trillion in 2004. To build successful e-business strategies, organizations will need to rapidly respond with secure electronic transactions over networks. Regardless of the access technology implemented for e-transactions, the networks will require to be secure.

Broadband networks based on digital subscriber line (xDSL), cable-modem, fixed wireless, and Ethernet-based metropolitan-area networks (MANs) refer to communication lines or services that operate at T1 rates (1.544 megabits per second [Mbps] and faster). xDSL, with two million consumers and one million business subscribers is projected to cover 70% of homes in the United States by 2002, while offering very low data and network security levels. Cable modems provide high-speed Internet data access to 30% businesses and an estimated 5.1 million consumers in the United States, while offering very low data and security levels. Fixed wireless is a developing market in 60 major metro cities, offering very low data and security levels. Ethernet-based metro networks are a developing market, providing services in 20 major metro cities and offering medium levels of data and network security.

Broadband networks share a common network cloud. The network cloud gets defined as each subscriber, business, or residence is added to the same segment and is managed by a common Internet service provider (ISP). These networks share a common network segment and are prone to threats related to sniffing of administrative passwords, data scanning, and transaction monitoring.

Today's threats are by-and-large the same as before, except that e-business has made it relatively easier to access the data and thus lowered the access threshold.

A recent computer crime and security survey of 538 security administrators from various computer industry, government, and academic organizations shows that the high price of computer crime and security breaches is on the rise. In 2000, the price of computer crime was estimated at $67 million, and in 2001 the price will top an estimated $151 million.

The Anatomy of a Network Exploit

A typical attack starts with establishing the footprint as the first step. As shown in *Figure 1*, this step is used to set up the scope of the attack through an analysis of publicly available information. The next step in performing a scan is to identify the vulnerability associated with each target of the attack resource. The step to follow is to exploit the target using a library of tools and techniques. This includes buffer-overflow exploits, address spoofs, and password and denial of service (DOS) attacks. The final step is to own the target causing damage, including blackmail, graffiti, information espionage, and destruction.

Computer threats and attacks can be broadly classified into two categories:

1. Random attacks
2. Targeted attacks

Random attacks include pervasive virus attacks, low-technology reconnaissance, port-scan attacks that exploit vulnerabilities inherent in the operating systems' software application configuration. In addition, the random hacks have been developed to exploit one or more vulnerabilities associated, for example, with integration of common gateway interface (CGI) scripts, hypertext transfer protocol (HTTP), or Java code. These attacks are more commonly orchestrated by the inexperienced script kiddie. In this case, the attacks are indiscriminate trolls on the Internet with little or no regard to who may own the networks.

Targeted attacks are performed by more experienced attackers who take time identifying the targets and conducting detailed reconnaissance before launching an attack. The attacks include DOS of prominent services, extraction and theft of vital information, or social engineering, physical break-ins, dumpster diving, war driving, and war dialing to locate modems.

DOS attacks allow the hacker to gain access to several computers connected to the Internet to execute a malicious code. The computer systems, at the hacker's signal, start sending data to the targeted Web sites.

FIGURE 1

The Anatomy of a Network Exploit

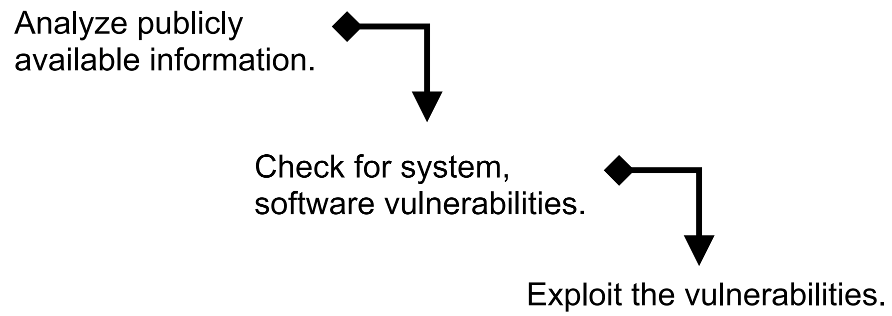

Social-engineering attacks are performed by an unknown person who calls on the telephone pretending to be another employee looking to verify sensitive information (including how to perform a particular task on a computer configuration) or inquiring about the company's stock performance, or pretending to be a system administrator calling an employee to fix his or her account on the system (which requires using his or her password). A successful means used by attackers to establish credibility in some organizations is providing a phone number that is on the organization's own phone system.

Attacks using dumpster-diving techniques are more effectively used for computer espionage amongst each other. Dumpster diving is a form of physical break-in. In this case, it involves rifling through an organization's garbage in an attempt to find sensitive information. As part of the search, the attacker may find a complete diagram of the network configuration or a carelessly tossed post-it note with an employee's log-in name and password.

War dialing searches for modem in a telephone private branch exchange (PBX) to obtaion access to a computer on the network. The war-dialing tool dials large pools of phone numbers to find unprotected network modems. As a variation, the demon-dialing tool searches for a password to gain network access to a known single telephone number.

Security of DSL and Cable Modems or Lack of…

DSL is an always-on technology providing a risky open door to the Internet. Attackers frequently scan the ISP address spaces that offer DSL services to exploit the computers of telecommuters and home-office workers. Although data may be shared by the same organization, with regional offices in different geographical locations, DSL provides network stovepipe connectivity and lacks a means to bind DSL sites securely, as shown in *Figure 2*. This exposes each Internet DSL connection, as any data traffic between sites is in the clear across the Internet. This risk is augmented by the fact that there are very few direct solutions serviced by several regional access providers that tie DSL back into the intranet.

In the recent distributed DOS (DDOS) attacks, zombies (or component codes) of the DOS attacks are being installed on vulnerable servers at universities or home-user systems connected to the always-on DSL. The attacker establishes groups of zombie-installed machines. These machines are in turn controlled by special control tools installed on additional exploited machines. The control tools instruct the zombies to conduct a DOS attack.

While secure tunnel-based authentication, session-based deployments are being used for xDSL–based systems, common point-to-point (PTP) clear text log-in and password data exchanges are more common. Owing to this shortcoming, many enterprise data information technology (IT) managers prefer to stick with the simpler-to-deploy dial-up solution.

Cable modem is a shared media network with an always-on connection analogous to the network segment of an Ethernet network. The shared media network is a combination of cable network subscribers for home cable entertainment, telecommuters with high-speed Internet connections, and small office/home office (SOHO) businesses.

A curious neighbor or malicious attacker, with sufficient knowledge of data-packet analyzers can access the data being exchanged by all subscribers of the network segment (see *Figure 3*). The network bandwidth offered by cable-modem networks is also temporary, and data access speed diminishes as the number of subscribers on the data segment increases. All cable subscribers share the same data pipe.

Security of Wireless Networks and Metro Networks or Lack of…

802.11b–based networks have rapidly gained popularity with the recent development of interoperability standards. The Wireless Ethernet Compatibility Alliance (WECA) can be accredited for the popularity of the 802.11b standards for interoperability, namely the wireless fidelity (WiFi) standard. Most vendors provide interoperability to the WiFi standard for data access points and network interface cards (NICs).

Fixed wireless networks use a local multipoint distribution system (LMDS) that decreases the deployment time of local access business circuits in the MAN. While the multichannel multipoint distribution dervice (MMDS) expands broad-

FIGURE 2

xDSL—Secure for Enterprise Network Access?

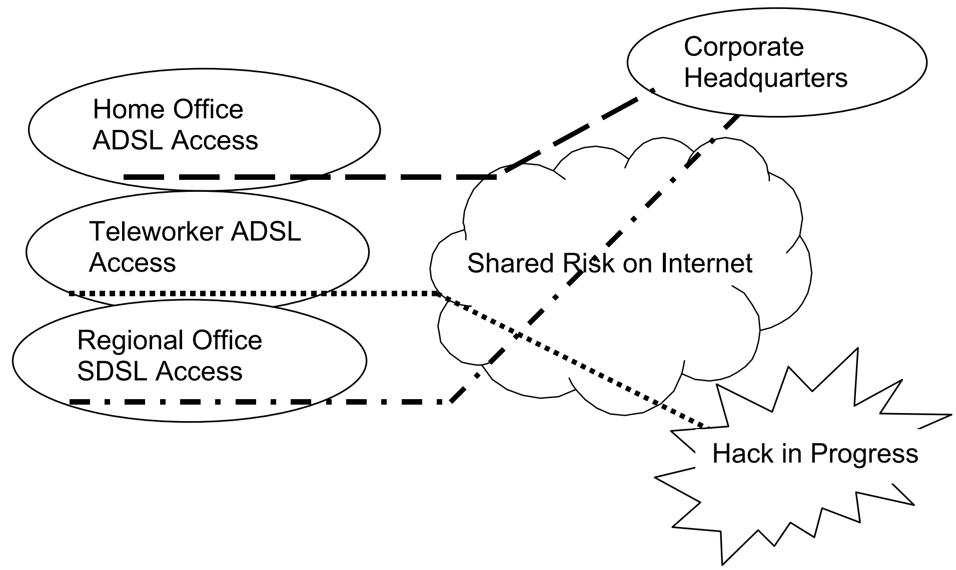

band access to consumers needing high-speed data access, the data is passed in the clear over the open airways.

The 802.11b standard includes Wireless Equivalent Privacy (WEP) provision for encryption. The WEP, depending on the vendor, offers two data-protection levels. One level is the virtually insecure combination of the 40-bit encryption key and 24-bit initialization vector (also called the 64-bit WEP encryption). The other level is the 104-bit key and 24-bit initialization vector (also called the 128-bit WEP encryption). Recent advances in the decryption of the RC-4 algorithm, owing to the key weakness, have been published and outline a method for extracting the master WEP key that will allow a potential attacker to pose as a legitimate user of an enterprise network.

Although the 802.11b standard employs weak static keys as part of the data encryption, there are an alarming number of access points to these wireless networks that are being deployed without WEP being activated. Furthermore, many wireless routers that provide data access have their default passwords. This data is captured by free data-analysis tools

FIGURE 3

Shared Medium Results in Shared Security Risk

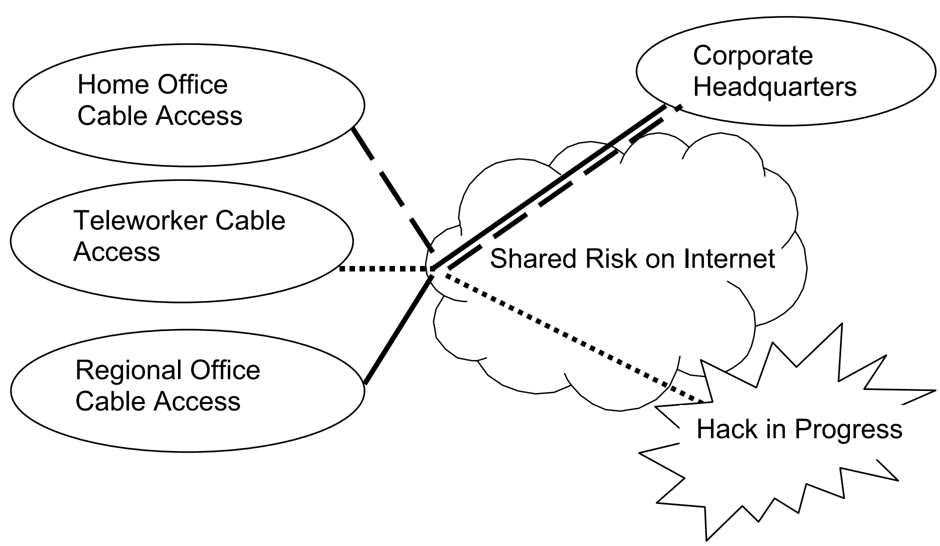

that are able to log vital system parameters, including the media access control (MAC) address of the access point, the name of the network, session ID (SSID), vendor name, data channel, WEP enable status, etc. Furthermore, a global positioning system (GPS) used in conjunction with the analyzer tool provides latitude and longitude for the data access point. Armed with the analysis, network information, and default passwords, malicious attackers can not only surf the Internet using the organization's network, but also change the router configuration and modify the Internet protocol (IP) addresses within the intranet.

Wireless network security is analogous to physical security. A malicious attacker given sufficient time, interest, and resources can gain access to the network. As shown in *Figure 4*, the malicious attacker with access to a wireless protocol analyzer, a Linux system with a 802.11b NIC, can listen to wireless data traffic communications between the basic local/remote service sets, which ase bridges between the wired Ethernet local-area networks (LANs) and the Internet.

The Ethernet MAN) offer a high-speed local wide-area network (WAN) connection. Limited by connection distance, MANs can extend from 70 km to 150 km. The virtual LANs provide secure high-speed links in the metro ring similar to the frame-relay permanent virtual circuits (PVCs). However, the long-haul connectivity back to the main office is over an insecure connection.

Security Alternatives

Broadband network-based data access using the PTP tunneling protocol (PPTP) connectivity provides a medium level of data security. This is because of rather weak encryption offered by the challenge-handshake authentication protocol (CHAP) and the remote authentication dial-in user service (RADIUS) embedded in Windows 95, 98, and 2000 operating systems. Data access using the Layer-2 tunneling protocol (L2TP) connectivity provides low-level data security using data encapsulation embedded in Windows 2000. Data access using the IPSec, site-to-site virtual private network (VPN) connectivity provides high-level data security. This provides data security using strong encryption (3DES), data authentication (HMAC and SHA-1), and user authentication (RADIUS and PKI) embedded in Windows 2000.

Routed private networks consolidate broadband pipes into a unified intranet (see *Figure 5*). A virtual private managed service network allows a subscriber to gain access to data through set of partitioned virtual routers (VRs). Each VR is the equivalent of an independent hardware router and is a gateway to the Internet. The virtual router is the backbone of the WAN with the firewall layered into the network. The layered approach allows secured sharing of networked resources for VRs to connect frame-relay and IP networks.

The VR enables customized services per subscriber, and data traffic is completely partitioned over the shared infrastructure. The subscriber data intermingles at the ISP VR just as it would in the ISPs' edge router.

Virtual firewalls or distributed firewalls provide a network-wide security with global policies distribution. This allows administrators to establish separate intranet and extranet policies per data site.

The trend in the industry is to deploy firewalls on individual laptop and desktop computers that access the Internet using

FIGURE 4

802.11b Network Exploit

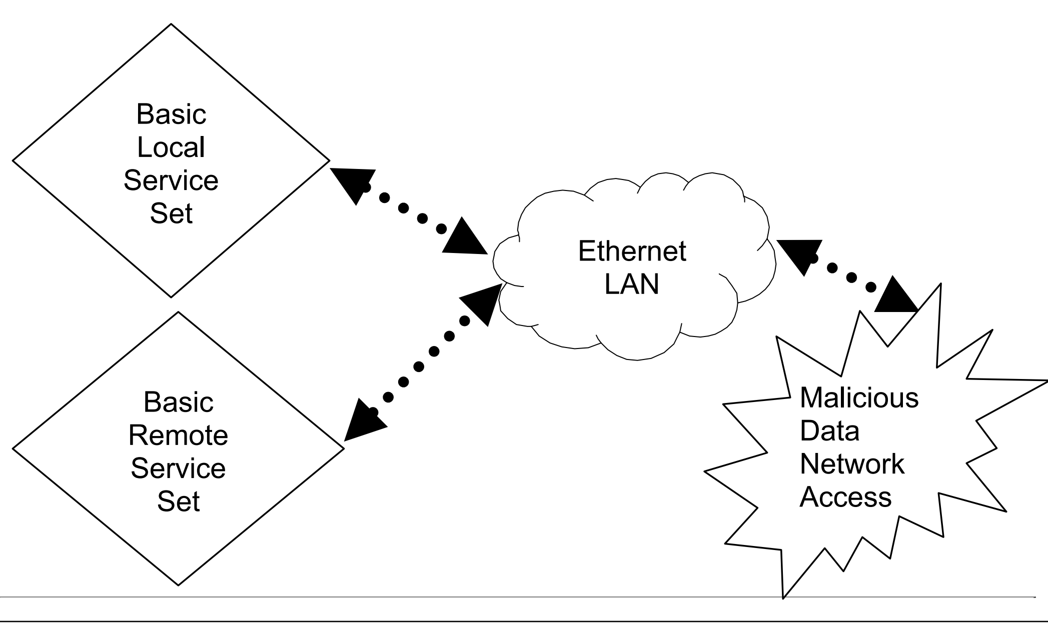

always-on DSL or cable-modem technology. These personal firewalls are deployed along with the corporate firewalls to providing additional data security. Malicious attackers scanning users with DSL and cable-modem services are easily blocked using personal firewalls. Working in the same way as the network firewall, the personal firewall is focused on monitoring network traffic coming into and leaving the computer. When the personal firewall detects malicious data traffic, it logs the traffic and blocks it prior to its causing potential damage. However, these firewalls are limited to monitoring the data traffic to and from the computer and do not analyze the installed programs or the operating system on the computer.

Traditional packet filters focus on monitoring individual data-packet header information and the data-packet direction. The data packets are monitored by defining a series of packet-filtering rules, also referred to as access control lists (ACLs), when they are defined on routers. Stateful packet filters extend the traditional packet-filter capability by remembering the data packets that have traversed the network. The defined rules control the decisions to either allow or reject data packets.

Wireless networks can be secured by ensuring that WEP is enabled. Next, it is recommended that the default SSID be changed, as the default ID is set using the manufacturer's default ID and password. Furthermore, the SSID should not be broadcast and should not point to the company address, as this will provide a relatively easy scan location. It is always recommended to change the default passwords for the access point or wireless router. Most network-analysis programs identify the manufacturer based on the MAC address.

More Secure than Passwords

Most users in an enterprise log into more than one system today to access e-mails and databases. Most of these systems require secure access with more than one user name and password. Each system has an individual account requiring its own user name and password combination. Users of these systems are faced with the daily challenge of dealing with several user name and password combinations, which in turn results in the exposing of vulnerabilities, such as when users write their names and passwords on Post-It notes sitting next to their desktop computers or monitors. Furthermore, some users have to share their access passwords with other individuals via an e-mail for system-maintenance purposes.

While single–sign-on systems can alleviate password-related security risks, enterprise-wide implementation of the single–sign-on is both expensive and complicated to deploy. Password-based access systems pose potential security risks to the enterprise owing to passwords compromises that may be accomplished using keystroke-logging, shoulder-surfing, and social-engineering techniques. In the event of a malicious activity using a password, the authorized user of a password does not know if his or her password has been compromised until the act is completed, because nothing physical has been taken. These risks can be minimized using the two-factor authentication hardware or memory tokens.

The intelligent hardware (disk-on-key) or memory-based tokens function as two-factor authentication devices and provide twice the assurance of a commonly used single

FIGURE 5

Secure Network-Based VPN Data Access

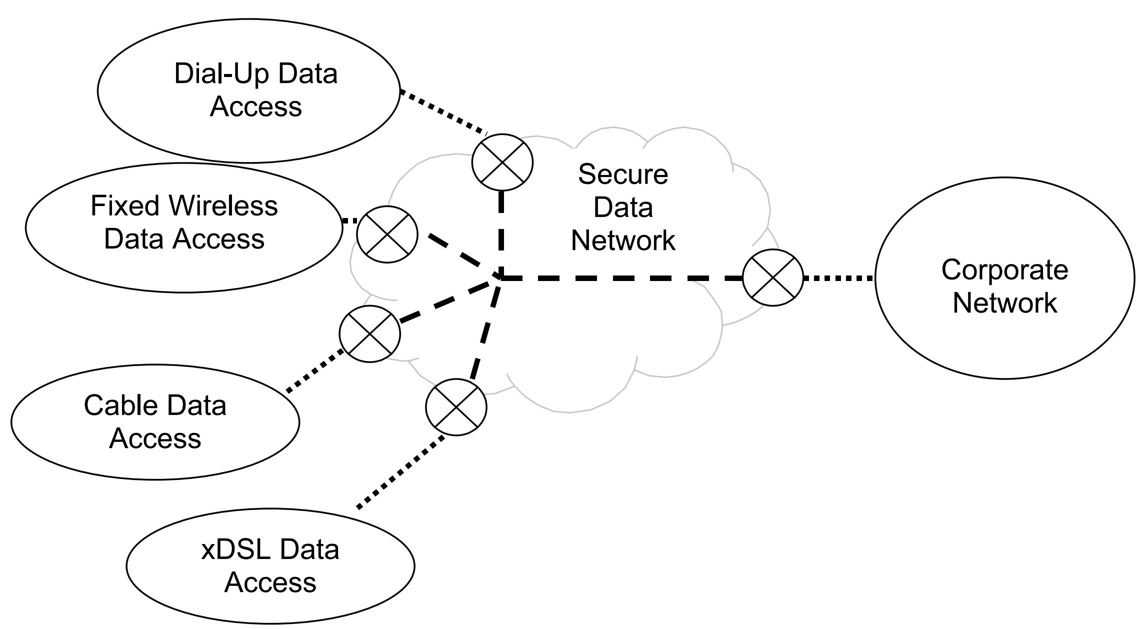

password scheme. While desktop or laptop password-based access is both easy to steal or compromise, the easy-to-use memory-based token or disk-on-key protects the shared secret, encrypted password, or the public/private key pairs by denying access to the laptop or desktop once it is removed from the universal serial bus (USB) or serial/parallel port. The secured public/private key pairs allow the authorized user to perform cryptographic functions, key generation, key negotiation, and user authentication. Only an authorized user, with physical access to the disk-on-key or memory-based token and the password can unlock the token.

References

1. *Securing Enterprise Broadband Networks*, Gartner Group Research and CoSine Communications Presentation (2001).
2. Ed Skoudis, *Counter Hack*. (Prentice Hall: 2001).
3. Craig Ellison, *Exploiting and Protecting 802.11b Wireless Networks* (2001).
4. Rob Schenk, Andrew Garcia, Russ Iwanchuk, *Wireless LAN Deployment and Security Basics* (2001).

Four Ways to Tell If Your Outsourcing Provider Is a True Business Partner or Just Another Vendor

Michael Filak

Senior Vice President, Outsourcing and Business Engineering

CGI

OK, you've done it. You've studied the global market and witnessed the shrinking windows of opportunity. You've identified three discontinuities: digitization, globalization, and deregulation. You realize how they are affecting your value chain and your supply chain. You've grasped how rapid technological change is making local companies global, disrupting value chains, and creating new competitors, many with fewer fixed assets and lower-cost structures than yours. You've also internalized a fundamental principle of life in e-Space: Information technology (IT) solutions must do more than simply create positive impact within your walls. They must also provide value throughout the extended enterprise—to your customers' customers and to your suppliers' suppliers.

Now you are ready to translate your worldview into action. In your search for the optimum IT solution, you have taken to heart objective voices of the academic community that say, "If an organization is not "best in world" at a critical activity, the organization is sacrificing competitive advantage by performing that activity internally or with its existing processes."[1]

You have studied industry analysts, such as G2R who say that "those companies who are best equipped for success in today's global information economy are not necessarily those with the largest volume of resources under ownership and management. Rather it is often their *access* to those assets, combined with traits such as speed of execution, intellectual capital and knowledge base, and business expertise—i.e., *specialization*—that result in work class performance. Such performance is achieved through innovative alliances with external providers."[2]

Yes, you've done it; you have made the strategic decision to outsource your IT. Further, you have identified qualified outsourcing vendors. All have favorable credentials in terms of financial strength and backlog. All will be around for some period of time. All can provide you the services that you need today, and all have the potential to provide you services in the future. Each of them promises to save you money on your IT. And while that is important, you also know that saving 10 percent to 20 percent on 3 percent to 5 percent of your operating budget is not the major driver for you to outsource.

Your primary outsourcing objective is this: You want to capture the opportunity side of IT—and that requires a true IT partner. The decisions you make now will affect the future of your organization. You are not ready to entrust that future to a "vendor." You want a partner who cares if you're successful—one dedicated to your success. You want a partner who sees outsourcing relationships as long-term commitments that must be win/win to survive. You want a partner who views quality IT services as the prerequisite to growing revenue, expanding profits, increasing speed-to-market, reducing competitive threats, and generating shareholder wealth. In fact, you know that financial strength, operating capabilities, and lower costs are merely table stakes for you to even consider a particular outsourcing provider as your strategic IT partner.

Test the Providers: Partner or Vendor?

The American Heritage Dictionary defines a partner as the following:

"One associated with another in an activity or a sphere of *common interest*."

The key point here is "common interest." Common interest means that you and your outsourcing provider are working together toward the same goals and objectives. Ideally, you also have a shared vision and a mutual passion for achieving the same positive outcomes for your business.

But how do you know that an outsourcing provider views *you* as a partner and not simply a cash cow? How can you be sure that you've entered into a relationship where through commitment, sharing of information, and intellectual capital a win/win partnership is established in which everyone moves toward the same goals?

To find out whether outsourcing providers view themselves as your true partners, try this simple test. Discuss the following concepts with each contender and see how they respond. Providers that embrace the following partnership models will aspire to a relationship deeper than a typical vendor/customer relationship. Those who don't, won't—and can be crossed off your list of partner candidates.

Flexible Business Model

First examine the provider's proposed business model. The model should make it clear that the provider is the delivery arm, and it must be flexible and adaptive to your needs, your industry, and your competitive marketplace. No "one size fits all" mentality here! For that reason, it is important that overall technology decisions and directions remain with you. Leverage the provider and the intellectual capital that is available. Seek out advice and recommendations. Even if you want the outsourcer to make all the decisions, understand that you ultimately determine whether or not you are successful and that it is your decision regarding how responsibilities should be divided. The chart in *Figure 1* demonstrates how responsibility is shared in a true partnership relationship.

As shown, you maintain what I like to call the "thin CIO" organization with ongoing responsibility for definition, authorization, supervision, and control. And, as it should be, the outsourcer appropriately becomes the delivery vehicle for you.

What's the flip side of this? Many outsourcers take the following approach: "Let us worry about technology; all you should really care about is the service." Outsourcers with this model are basically telling you that you need to adapt to the their business and technology model. The problem is that this model is created to achieve the greatest leverage for the provider. Leverage that must be gained and shared across many different industries and customers whose overall requirements are not the same as yours. Tread carefully with providers that approach you with this "least common denominator (LCD)" mentality. The technologies and models that best serve your needs may not fall within this LCD scenario. Additionally, should you ever decide to bring your IT processing back in-house; you are stuck with your provider's environment. This is even further exasperated when you choose an outsourcing vendor that is aligned with its own hardware or software. Technology decisions are made based on the vendor's product lines—not your needs. Letting your outsourcing provider make technology decisions may earn you an extra 3 percent to 5 percent savings off your operating budget. Yet squeezing out every last penny of savings is not the main driver of outsourcing in the first place. For this, you may be sacrificing true technological and competitive advantage.

No, a true partner puts your needs first. Setting technology direction is a collaborative process between you and your partner. However, the final decision is yours. You must feel confident that the technological solutions that you deploy are the technological solutions that drive success in your business.

Flexible Contracts

Outsourcing contracts are, by nature, longer term—typically seven to 10 years, although 15 years is no longer uncommon. There are very good reasons for this. Unlike typical

FIGURE 1

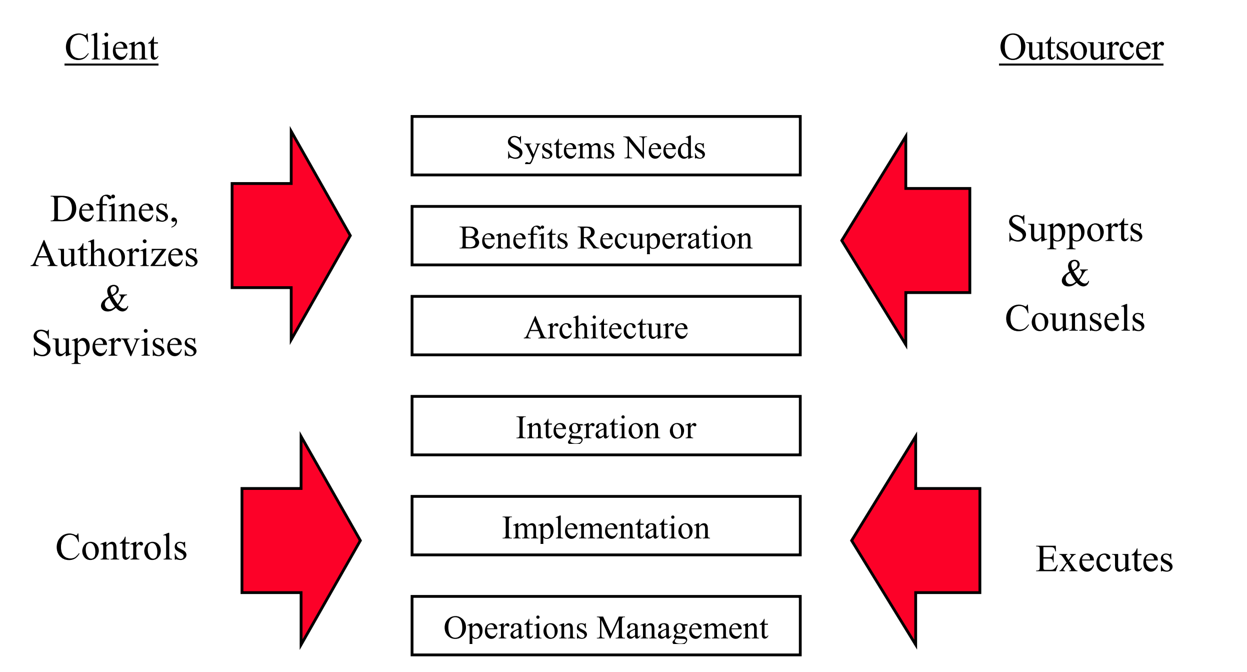

vendor relationships, outsourcing relationships closely resemble merger and acquisitions (M&A) in that assets and people are transferred into the provider's organization. Additionally, the provider commits to increased performance at a lower overall cost. The outsourcer's necessity to recoup the financial investment required to purchase assets, transfer people, and implement tools, methodologies, and processes to reduce costs typically takes several years.

However, as time passes, your business changes. Your company changes. Your competitors change. Your marketplace changes. Your strategy changes. Yes, the world goes on while your outsourcer tries to recoup its investment. This is when you may recall the warnings you have read about locked-in outsourcing contracts—a situation where your contract is ruling your business decisions. This is unacceptable and not reflective of a partner relationship.

A true partner with your best interests in mind will recognize that the contract must be a living, breathing, and changing document that captures the spirit of the relationship rather than the absolutes. Without you even needing to ask, a good partner will recommend a process or methodology for amending the contract to accommodate your changing business. After all, a partner knows that successful relationships must be win/win. And you recognize that a successful relationship must be win/win.

Your current contractual arrangement must flexibly respond to business changes—even if the result is lower revenue to the outsourcing provider. For example, say your company is going through a period of "refocusing," and you need to sell off non-core portions of your business. Yet your outsourcing contract ties you into certain services and a minimum-billing figure. Obviously, that contract no longer meets your needs. Together, satisfactory arrangements can be made if everyone is working toward a win/win relationship.

Learn up front how potential providers view contract changes. In fact, watch carefully how they negotiate the contract. Do they spend much time and energy protecting themselves from any adverse event no matter how minimal the chances that it may occur? Study any provider's track record. Do their words during negotiation carry over into future issues? It goes without saying that any provider that is the least bit rigid on mechanisms to change the contract cannot be considered your partner.

Open-Book Financial Model

Are prospective outsourcing providers open to sharing financial information directly related to your contract? Sharing of information, especially financial, is a true test of a partner. Providers that cannot take an open-book approach with you are not true partners. Now a typical response may revolve around the issue of benchmarking. Some providers may offer a benchmarking provision to lend an illusion of frank and open-business practices. The provider may say, "I won't share my financial information with you, but I will (reluctantly) let you utilize a third party to benchmark my pricing to you to ensure that I am charging you competitive rates." However, there is usually a caveat—"competitive for deals with similar size and scope." The major problem here is that typically there is no exact match on similar size, service, and scope. So the benchmarking provision may provide no value at all. Similarly, there are times where the rates may be market rates, but the provider is realizing enormous margins and not sharing the windfall with you.

No, a true partner has nothing to hide and should be more than willing to share the financials regarding services to you. This openness works to your benefit in two ways. First, by accessing the financials, you can be assured that you are paying a fair price for your services. After all, the provider's internal cost structure must be competitive to guarantee success in the marketplace. Second, you will be able to see your partner's margins. Nothing is more destructive to your outsourcing relationship than a growing feeling that your outsourcing partner is taking advantage of you. The sharing of financial data alleviates this issue. This openness fosters trust between the two parties. Should technological breakthroughs occur that have a significant positive impact to the outsourcer's cost structure, you will be able to see it. A true partner will be flexible toward changing the overall pricing structure to share these savings with you. The flexibility of your contract should ensure that you are not locked into a financial structure where you become a cash cow to the outsourcing provider.

Now, what is the best way to accomplish this open-book approach? I recommend that a third-party auditor be retained to review the books of the outsourcer. Ahead of time, both parties meet with the auditor and define the goals and objectives of the audit.

Another financial issue to look at closely is the manner in which the outsourcer leverages its technology alliances for you. These technology alliances should be built for the sole purpose of providing better and more cost-effective services to you. For instance, how does an IT outsourcer leverage its buying power on your behalf? If the outsourcer can purchase a particular piece of hardware for 25 percent less than you can, how is that passed on to you? A true partner passes on the total savings, less a small management fee, to its clients. Another method, though not quite as partnership-oriented, is to split the savings equally. However, the most common method amongst non-partner–oriented outsourcers is to plug the entire amount plus overhead plus profit margin into the price. In this case, the outsourcer functions as a reseller. Besides the financial implications, this method also removes objectivity from the outsourcer, which will look to increase their profit by recommending technology (hardware and software) where the outsourcer can realize the largest profit.

Transition/Migration Continuity

During the sales process, competing outsourcing vendors may parade before you a cadre of experts to impress you with their technological and business prowess. You need to ask each vendor, "How many of these experts will stay around and be directly involved after the contract is signed during the transition/migration process?" During the sales process and subsequent contract negotiations, an abundance of intellectual capital is generated as to how this contract will operate. It is extremely important that the contract gets off the ground running as per negotiations. The learning curve must be minimized. It is imperative that key personnel remain committed and involved throughout the transition/migration process. If the outsourcer's sales/pursuit team merely cedes the operations to another team, you can be assured that they are off and running, looking for their next cash cow. Once again, a true partner is just as interested (if not more interested) in making sure that your contract works to your satisfaction and benefit as they are in finding new contracts.

Careful Evaluation Now Yields Positive Results Later

Before you enter into a long-term agreement with any outsourcing provider, carefully investigate the four partnership concepts presented here: flexible business model, flexible contracts, open-book financial model, and transition/migration continuity. Do your prospective IT outsourcing partners readily accept these concepts? Do they have to be pulled "kicking and screaming" to your position, or, worse yet, do they flatly refuse? Remember that outsourcing agreements are more than just bottom-line revenue support of lower IT costs and better service. The best agreements focus on top-line revenue—generation of increased competitiveness, enhanced services, speed-to-market, expanding markets, and an overall internal focus on your core competencies. These vital goals can only be attained with a true partnership relationship with your IT outsourcing provider.

Notes

1. James Brian Quinn, professor emeritus, The Amos Tuck Graduate School of Management Dartmouth College; Winner of the 1999 Outsourcing World Achievement Award
2. *G2R*

Architecting Software Systems for Network Availability

Chris Ford

Senior Product Manager

Équipe Communications Corporation

Introduction

In the last few years, equipment vendors have been focused on developing products for carrier networks that meet the skyrocketing demand of data communications. Carrier networks are built and continuously expanded to satisfy the relentless expectation of a paper-free, instant-information society. Our dependence on networks continues to increase exponentially, placing further pressure on equipment vendors and carriers to build networks that are resilient to change, growth, and intermittent failures.

During the past few years, we have seen numerous catastrophic data-networking failures that have halted productivity for many sectors, including financial, education, and retail, at a cost of hundreds of millions of dollars. Several outages within Internet service provider (ISP) networks and at network access points (NAPs) have resulted in widespread network unavailability. Major frame-relay (FR) interexchange carrier (IXC) networks crashed and remained down for as long as 11 days, leaving thousands of high-paying corporate customers without service.

A recent study conducted by Sterling Research, Inc. estimates that an hour of downtime for a retail stock broker in New York City costs as much as $6.5 million and that a "down" automated teller machine costs about $15,000 per hour. However, it's not only the end customer who suffers from the network outages: Rick Malone, a principal at Vertical Systems Group, indicates that recent domestic IXC outages have cost those carriers more than $200 million. The service market for data communication is *billions* of dollars, which means that a single drop in market-share caused by network outage translates into tens of millions of dollars in lost revenue.

Perhaps more disturbing than the frequency of recent data-network failures is the anticipation of what the future holds. New technologies such as multiprotocol label switching (MPLS) and generalized MPLS (GMPLS) are emerging to allow carriers the freedom to converge voice, video, and data traffic—all across a common backbone with a single control plane. Service-level agreements (SLAs), virtual private networks (VPNs), customer network management (CNM), and other value-added services are also adding tremendous complexity to the relatively simple networks of today. As carriers migrate toward single, collapsed data networks and advanced data services, the stability of the public network will be increasingly challenged. Thus, solving the network reliability problem is an imperative. Given the inevitable increase in size and network complexity (see *Figure 1*) and the fact that networks are still running on software architectures designed in the 1980s, the problem is poised to get worse, not better.

The underlying cause of a vast majority of network failures has been *software*. Management synchronization issues, runaway software faults, memory leaks, data corruption, inadequate upgrade procedures, and routing storms that flood networks with control traffic are all examples of how software can impact network availability. Unfortunately, software reliability has been neglected in favor of meeting demands for new features and for speed to market in delivering them. As a result, vendors in the data communications field have failed to introduce new, more reliable software architectures.

This paper will identify the following:

1. Software architectures that exist today and why they continue to make networks vulnerable to outage
2. A strategy for designing software architectures that will provide true carrier-class reliability for next-generation networks

Limitations of Today's Legacy Software Systems

This overview provides the opportunity to shed light on some legacy networking platforms, underscoring the need not only for a new class of networking equipment, but also for a new class of networking software.

Let's start out by analyzing the anatomy of a typical switch/router within the network. The trend in hardware design is toward modularity, flexibility, and scalable, nonintrusive growth. These very principles can and should be applied to system software. For example:

Carrier-Class Hardware: Today's switch/router is rack-mountable with many slots in which a variety of hardware modules may be inserted, hot swapped, and upgraded. Modular hardware enables carriers to do the following:

- Scale as required (pay as you grow)
- Hot swap modules without bringing down the switch/router
- Purchase media-specific modules

Today's modular hardware provides flexible "speeds and feeds" and reasonable reliability required for network growth. This design approach has become more common within the industry.

Legacy Software: Today's switch/router software (see *Figure 2*) is still based on a single, monolithic image that is typically bound to the hardware. This can make the software:

- Vulnerable to a single fault (A rogue pointer or memory corruption can bring the switch/router to its knees, forcing a restart.)
- Fragile under network load (An excessive load can force a system panic and automatic reboot.)
- Difficult on which to perform problem isolation and pinpoint the root cause of software problems
- Difficult to upgrade, since extensive testing cycles are required to validate each release
- Unwieldy and difficult to manage as it grows with every release

In fact, a monolithic code base can drastically increase the amount of time it takes to isolate and control faults, pinpoint the root cause of the problem, and deploy a fix. Not only does this reduce the overall availability of a service-provider network, but it also impedes the deployment of new features and services. This is evident in the life cycle of system software, depicted in *Figure 2*.

Today's monolithic software does not meet the scalability or reliability requirements of highly available networks. Monolithic images are becoming increasingly problematic and unable to meet the needs of tomorrow's networks.

The solution is to leverage the principles of hardware modularity and advances in software technology to achieve a software solution that scales linearly to meet tomorrow's growth and reliability needs. Next-generation hardware requires platform software that stays synchronized with the element manager even during failure. It also requires that faults be contained according to predefined policies, without system reboots. Lastly, platform software must allow for upgrades that are non-disruptive to customers and result in predictable outcomes. Scalability, reliability, and availability: the cornerstones of the carrier-class platform software architecture.

Leveraging Distributed Architecture

Within networks today, distributed architecture is widely used to provide performance, scalability, and overall network availability. The Internet itself can be viewed as a distributed system comprised of thousands of loosely coupled elements that work together to provide the end user a single, richly connected network.

A well-constructed software design can be viewed as a network of loosely distributed processes that work together,

FIGURE 1

Network Complexity and Size over Time

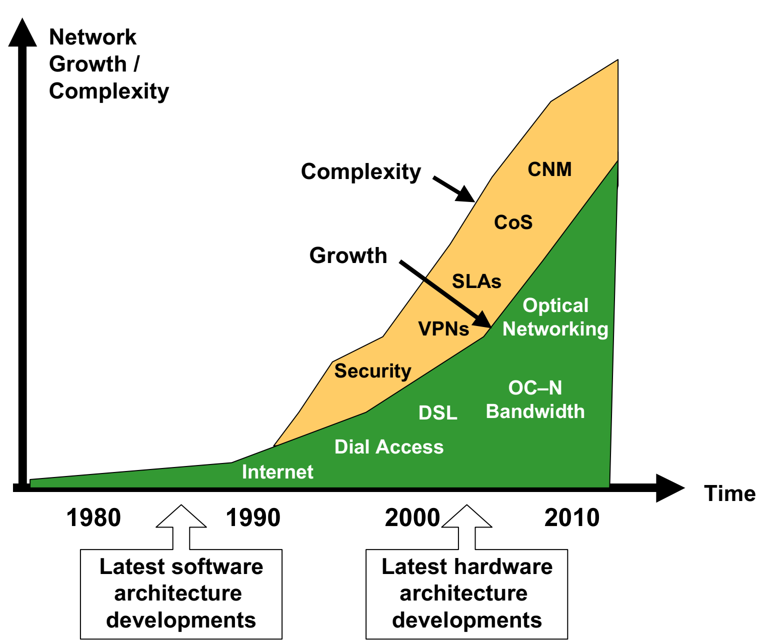

FIGURE 2

The Life Cycle of Legacy System Software

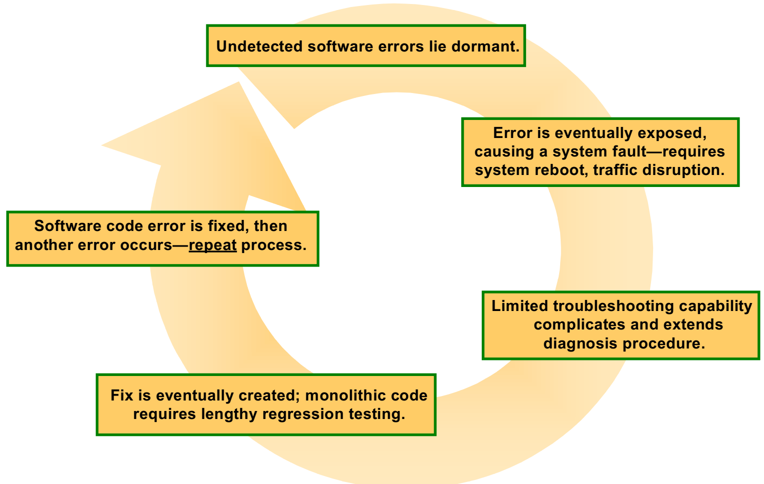

not unlike the processes that comprise the Internet. Like the Internet, it leverages a distributed architecture to achieve the performance gains of parallel processing to provide system scalability and to deliver better overall availability of network service.

Distributed Control Architecture

Many legacy switch/routers have built on the aforementioned theory, adopting distributed architectures within the data plane to achieve the aggregate bandwidth required for today's networks. By running many "fast forwarding engines" in parallel, data can be switched or routed at higher rates of speed, increasing the overall forwarding performance and subsequently enabling networks to scale.

While distributed processing techniques are currently used within data planes, the control plane has remained largely centralized. Until now, this architecture has performed reasonably well and was simple to program. However, many of the same forces that shaped the evolution of the data plane are now exerting similar pressures on the control plane. The next generation of switch/routers will run resource-intensive control protocols that will be difficult to scale in a centralized architecture. Some of the features to be supported include the following:

- Signaling protocols
- Connection path routing
- Differentiated services
- Software interface scaling

Because these control services are hungry for central processing unit (CPU) cycles and system memory, a *distributed control architecture* is required to scale next-generation platforms and provide the desired level of network availability (see *Figure 3*). Fortunately, these services are relatively straightforward to distribute: They are localized to individual interfaces and require little, if any, shared global state. When these applications are deployed in a distributed fashion, their performance and scalability will increase linearly as processing modules are added to the platform—far exceeding the performance and scalability that can be achieved in any centralized control architecture. Distributed control architectures are a prerequisite to achieve the scalability required in carrier networks.

Location Transparency Provides Significant Flexibility

To offer the flexibility needed for today's networks, application software must be constructed without concern for the location of other services on which it depends. In other words, software components must have *location transparency*. With location transparency, software designers can focus on clean design partitioning, independent of the hardware architecture. This approach insulates the software system from changes in hardware as the product evolves.

In addition, a flexible distribution scheme enables software components to be relocated, tailored, and optimized without redesigning the core system. This flexibility is especially important, since the design team often learns more about the processing, communication, and resource-consumption characteristics of a product as it is tested and used in production networks. The result is a scalable, flexible software architecture that can easily be adapted to support emerging product requirements, new product families, and even new hardware platforms.

Improved Availability

Perhaps the most significant benefit of adopting a distributed control model is the impact on availability. When a component fails in a centralized architecture, it has an adverse effect on the entire system. A significant loss of data, control, and management services usually occurs. In contrast, failures in a distributed architecture are localized to the affected component. Only a small fraction of the system's services is affected. The outage is significantly shorter because there is much less work required to recover. A good distributed control software must mitigate the effect of single component failure by isolating the failure to the offending component. Service disruption is contained to services

FIGURE 3

Distributed Control Architecture

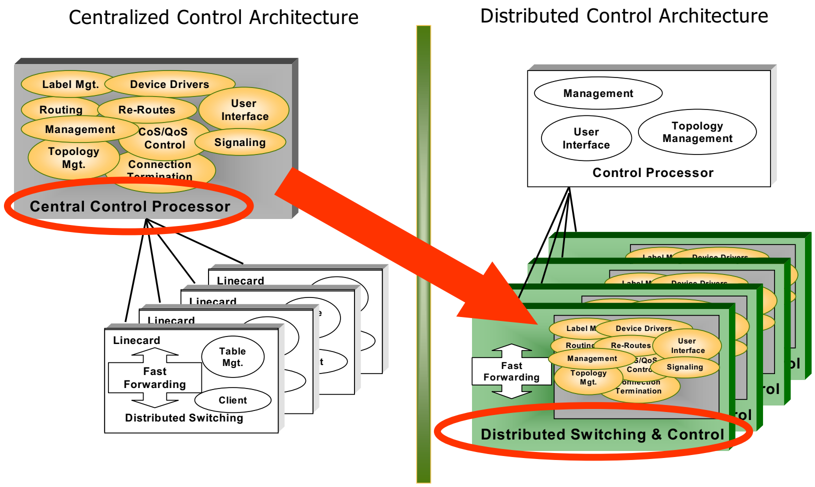

that utilize the offending component rather than all services provided by the entire system.

Automated Fault Detection and Recovery

Even the best, most reliable software system will contain defects that generate system faults. The key then is to properly manage these faults so that system disruption is minimized. Thus, fault isolation and containment are fundamental to any next-generation software architecture. Several developments in operating system (OS) technology enable next-generation platform software to reduce the likelihood of software faults and several mitigate and contain the effect of software errors.

Use of Protected Memory

Most contemporary switch/routers are built using a real-time, monolithic software architecture. These architectures achieve efficiency by running all programs in unprotected kernel mode. If a failure occurs in any piece of software (e.g., a corrupted pointer), the entire system will crash or become unstable. Instability is often worse than a hard failure because system operation can degrade severely without any indication of the cause. The only way to recover from a failure in these systems is to reboot. Unprotected, kernel-mode architectures make it extremely difficult to achieve high-availability goals.

A higher level of reliability may be realized by using hardware and OS features to implement a *protected memory* model. In this model, each application runs in its own protected memory space, ensuring that the failure of one application will not adversely affect the operation of others. In the earlier example, an application with a corrupted pointer will terminate before it can do damage to any other part of the system. Employing a protected memory scheme eliminates much of the instability and indeterminism seen in monolithic architectures. As a side benefit, protecting applications also facilitates the development of clean, well-defined application programming interfaces (APIs), a prerequisite to the software modularity discussed earlier in this document.

The Importance of a Microkernel

A few newer switch/routing platforms have made the shift to a programming model that uses protected memory. However, many of these are built on top of large, UNIX–derivative kernels that still compromise the reliability and modularity of the system. In these systems, there is a significant number of OS components and system device drivers that run unprotected in kernel mode. A failure in one of these unprotected components will cause a reboot of the system. In addition, time-critical control tasks are often written to run in kernel mode (as shown in gray in *Figure 4*) to avoid the performance penalties of a large OS. As the amount of code running in the kernel increases, the likelihood of a kernel fault rises dramatically. A microkernel architecture greatly improves system availability. Only the most basic system services (e.g., interprocess communication and scheduling) should run in kernel mode. All applications, device drivers, and higher-order system services run as user processes. This approach greatly reduces the amount of code running in kernel mode. The following characteristics are important qualities of any microkernel architecture:

- *Reduced Size:* An efficient microkernel, an order of magnitude smaller than a traditional kernel, greatly reduces the possibility of a service affecting failure.
- *Granular Process Control:* Processes can be started and stopped independently. The microkernel tracks all process resources so they can be recovered in the event of a failure.
- *Superior Performance and Scaling:* High-performance and scalable applications can be developed without moving the applications into kernel mode.

Traditionally, fault detection and monitoring do not receive a great deal of attention from network equipment designers.

FIGURE 4

Full-Protection, Microkernel Architecture

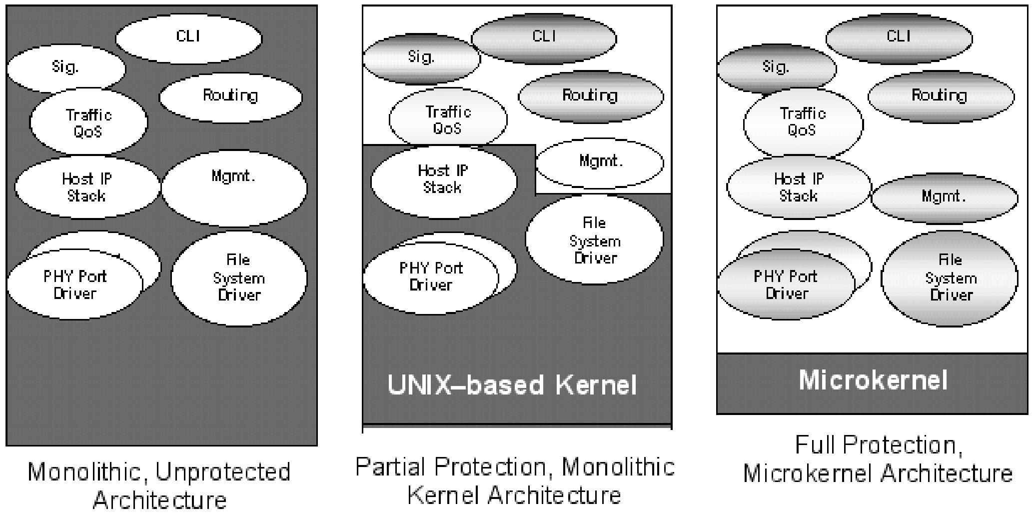

To reach high availability goals, however, systems must adopt a proactive approach to fault monitoring. To provide comprehensive fault detection, switch/router-resident management software must intelligently monitor components and determine when action needs to be taken. It must do the following:

- Unobtrusively inject heartbeats throughout the system
- Query active components for status in real time
- Monitor the resource (CPU, memory, etc.) consumption of each software module
- Actively monitor redundant hardware and software components to ensure that they can take over successfully in the event of a failure

This strategy provides coverage for a wide array of failure conditions and automates the detection and notification processes. The result is an intelligent, cohesive, automated fault-monitoring service.

Adaptive Fault Handling

Part of effective fault handling is customization and intelligent processing of faults. To be adaptive at handling faults, system software must do the following:

- Use an intelligent, policy-based recovery scheme (e.g., recovery policy may dictate process restart, process re-initialization, module switch over, process termination, or another corrective action)
- Maintain a complete history of system faults
- Leverage heuristics to determine whether a failure is indicative of a more severe problem and automatically escalate the recovery action

Rapid Fault Recovery

Speed of recovery is one of the primary considerations for picking an appropriate fault-recovery action. Restarting a single software process running on a software module is much faster than switching an entire slot to a redundant component. Clearly, when a single process is restarted, only a fraction of a module's services is affected.

Two measures may be employed to ensure rapid recovery from hardware and software failures. First, the protocol stacks may be multithreaded, so the failure of a software process affects only a fraction of the module's capabilities. For example, an eight-port module might have eight separate instances of the protocol stack. If one process fails, only one port's service is affected. Second, all software entities are designed to recover from failures of an adjacent or peer entity by performing dynamic state *checkpointing*.

Checkpointing

For a backup component to take over from a failed or removed component without operational downtime, the backup needs to receive and maintain current state information from the active component. This transfer of dynamic state is called checkpointing. With checkpointing, the following are done:

- The backup process state is maintained by sending information from the active to backup components at regular intervals.
- Dynamic state can be used to recreate a module's processes on the designated redundant module in the event of a failure.
- Recovery from a software failure is normally complete in milliseconds, resulting in no discernible service disruption. With an entire hardware module failure, it may take only a few seconds to restore service—a great improvement over existing solutions.

More important, the affected interfaces never "go down" from an operational perspective, so the routing topology and forwarding tables remain intact. There is no need for the other modules in the system (and other network elements in the network) to reroute around the temporary failure.

Unprecedented Availability

Performance under stress, day-to-day management, and smooth periodic upgrades are important aspects of the network life cycle. Network integrity must not be compromised with normal or abnormal traffic spikes or unplanned network events. In addition, software modularity, innovative management techniques, and upgrade tools are an important part of providing unprecedented software availability.

Stability under Network Load

One common and disastrous problem seen in data networks today is that of instability under severe network loads. Often, this network load is created by control traffic. It is for this reason that a discrete separation must exist between the control and data planes. Using a dedicated control plane enables the distribution of high-performance software applications without affecting the data plane of the platform. Priority queuing on the control plane ensures that critical control traffic (e.g., routing protocol updates and "keepalives") receives priority service during times of heavy network congestion. This will ensure that excessive control traffic does not interfere with user data and that dropped control traffic does not further exacerbate the amount of congestion.

Software Modularity

Since modularity is a fundamental technique for controlling complexity, its application to a fault-tolerant, distributed system brings significant rewards. Software modularity is the practice of functionally dividing a software system into individual modules, which can then be designed and implemented independently and later integrated into a cohesive solution. All interactions between modules are made through well-defined APIs, ensuring that modules can continue to be developed independently.

Modularity is a powerful tool for the following:

- Controlling and simplifying the design, construction, testing, and maintenance of complex software systems
- Encouraging the reuse of functionality, which in turn improves efficiency and minimizes bugs
- Facilitating customization and extensibility of products
- Enabling less intrusive, more contained software upgrades

Discipline and enforcement are the keys to delivering true modularity. Core system services can encourage disciplined treatment of APIs by providing one (and only one) way of communicating between programs. This forces developers to communicate through the main API because it's the only way to exchange information between two applications. The creation of back-door mechanisms, which compromise software integrity, becomes impossible.

Because of this strict enforcement of APIs in software modularity, the use of shared memory between processes is forbidden; *message passing* is the single interprocess communication method. Protected memory helps to enforce this physical separation of software modules. Every logical component is compiled and linked as a separate program, and each program runs in its own protected memory space. There is no way to share memory between two separate programs, so the programs are forced to communicate via message passing. Thus, the software system's modularity also helps to minimize the disruption associated with deploying a new software release. As will be seen subsequently, individual software modules (fixes or features) may be upgraded without impact to the rest of the system.

Release Management

Suppliers of monolithic software are forced to adopt a very conservative release strategy because of the logistical problems associated with testing and deploying large, single-image software. Typically, major feature releases can be delivered just once or twice a year, and "patch" releases must target a large collection of customer concerns. This conservative strategy limits the vendor's ability to react to evolving customer needs. Given the rate of change in the networking industry, this approach can severely hamper a provider's ability to upgrade the network or offer new services.

Modular software design allows a platform vendor to adopt a progressive release management strategy with several benefits:

- New software modules and device drivers can be released to the field as soon as they are available.
- Fixes and features can be safely applied to a subset of elements within a network, significantly reducing the number of wholesale network upgrades.
- Individual components of a modular system can be tested in isolation with confidence, facilitating higher-quality, timelier releases.

Upgrade/Downgrade Agent Features

In terms of system availability, operator errors during software upgrades can be a significant contributor to system downtime. To minimize these errors, upgrades to software must be automated. An intelligent "upgrade agent" would be charged with determining the scope of the upgrade and applying the necessary changes to the affected modules.

Manageability

No comprehensive system design would be complete without careful attention to manageability. The platform must leverage a configuration repository that provides a centralized, standardized, and flexible way to access management data. The repository must provide a set of APIs to enable applications or other parts of the management framework to read, write, and search for data. Rather than using simple network management protocol (SNMP) for this information repository, vendors should consider a transaction-oriented approach, which makes use of relational database technology embedded in platform software. Database transactions can ensure configuration integrity, eliminating partial, incomplete (and thus invalid) configurations. Other important features include recoverability and replication of configuration data. Relational database technology provides a rich suite of features that go far beyond traditional SNMP–based solutions.

To eliminate synchronization issues with an element-management system (EMS), the relational database in system software pushes configuration change to the administrative system, eliminating the need for excessive polls out to the device.

Summary

For networks to achieve the scale, reliability, and availability demands placed on them, a new software architecture must accompany next-generation network hardware. It must be distributed, modular, and include integrated proactive fault detection, isolation, and correction. This revolutionary approach must be built from the ground up—adding scalability and fault tolerance to a legacy software architecture is all but impossible.

Network Optimization Challenges

Daniel Gatti
Consultant
Mayan Networks

This paper focuses on network optimization challenges and their many contributing factors.

Network optimization is big business. Within the past two years, more than 100 vendors developing new solutions and products all geared at optimizing parts of the network have been funded. Within the past six months, a number of companies have been acquired and a few have gone public.

A Question of Fiber

Obviously the network means different things to vendors and service providers. There are 39 million miles of fiber in the United States today, representing tremendous capacity in the long-haul network. In the last two years, the metro core has also been built out with a significant amount of fiber. At the edge of the network, fiber is scarce, and the prevailing physical medium continues to be copper.

There are approximately 770,000 class-A commercial buildings in the United States today. Only seven percent of those building have fiber connectivity. When evaluating network optimization strategies, one of the key issues that remain to be solved is how to get fiber to those major customers at a reasonable cost. Several new carriers are bringing fiber directly to these customers but are hampered by prohibitively high costs associated with carrier hotels and digging up streets.

Optimizing the network to take advantage of the existing infrastructure is a huge challenge. In the United States today, there are 150,000 synchronous optical network (SONET) rings in place. The rest of the world has another 150,000 synchronous digital hierarchy (SDH) rings. The challenge is to take advantage of this capability, try to use it, and maximize the fiber capacity.

Metro Optical Market

Traditional SONET equipment, including basic add/drop multiplexers, accounts for a constant $4 billion-per-year market. However, according to Pioneer Consulting, the metro optical market will grow from less than $5 billion in 2000 to $14.3 billion by 2004. The fastest growth in the metro optical area is in multiservice SONET, or multiservice provisioning platforms. Pioneer expects multiservice SONET to grow from less than $500 million to $6.3 billion by 2004. Multiservice SONET combines traditional time division multiplexing (TDM) functionality with asynchronous transfer mode (ATM), frame relay, Ethernet, or Internet protocol (IP). Optical IP is the fastest-growing service, increasing from relatively nothing to $2.5 billion by 2004. Therefore, a considerable amount of optical capability is coming into the metro.

Higher-layer protocols and services continue to grow. Current thinking has it that packets will dominate traffic volume by 2003, but that multiservice traffic will also grow. Frame relay and ATM growth in the last quarter of 2000 was surprisingly high, and even though the move is toward IP and higher-speed Ethernet, traditional ATM and frame relay revenues continue to grow for carriers.

The Problem of Multiple Protocols

Companies can optimize their networks if they attack the issues at the edge. First of all, multiple protocols make provisioning of services very difficult. A combination of services exists today—ATM, frame relay, voice, TDM, and obviously IP are coming in—but the challenge is how to provision those services to the end-use customer.

Some buildings contract with several different service providers because bundled service is not available from any single one of them. One provider may supply T1, while voice comes from another, digital subscriber lines (DSL) from another, and Internet access from yet another. However, customers do not want to contract with multiple providers, preferring bundled services at an attractive price.

The Problem of Multiple Devices

With each additional protocol that a service provider deploys comes a separate piece of equipment that requires installation, configuration, and maintenance. Whether frame relay devices, ATM devices, or integrated access devices, all require some physical change of equipment and a truck roll—an extremely labor-intensive process. Today's SONET/SDH-based infrastructure was never built to handle data traffic, and a 20 percent utilization of existing SONET rings is the best companies can expect when deploying data services. The inefficiencies of data transport over SONET, configuration and maintenance complexities, and the cost of a truck roll are the key challenges in network optimization.

Dealing with the Network Edge

All of the major carriers have built a huge amount of capacity in the core of the network. Raw fiber capacity in the core creates a real bottleneck at the network edge. The challenges are to bring that core capacity out to the edge, to aggregate traffic and services efficiently, and to transport them back into the core. These disparate issues make edge optimization very difficult.

The complexity at the network edge comes largely from the variety of interfaces and protocols that must be supported. One building may have private branch exchanges (PBX), DSL, and customer-premises equipment (CPE). That equipment includes routers, Ethernet switches, and all kinds of integrated devices handling not only voice, but virtual private networks, storage-area networks, and web hosting. Multiple physical interfaces—including Ethernet, T1, T3, digital subscriber line (DSL), optical carrier (OC)-3, OC-12, and OC-48—are coming in with different protocols, whether TDM, ATM, frame relay, or IP. In addition, new protocols and approaches are constantly being developed. Provisioning this diverse mix of interfaces and protocols is extremely labor intensive and expensive.

The Metro Edge Is Key

The challenge of optimization revolves around service delivery at the metro edge. Service providers must deploy both processes and equipment that reduce the cost and complexity of turning up new customers, bring core bandwidth to the network edge, and increase revenues through reduced operations costs. Those carriers able to achieve this level of network optimization will enjoy great success in the years to come.

A Step in the Right Direction toward Accurate, Consistent, and Reliable Information

Tom Glover
Customer Unit Leader, Communications Industry Sector
Unitech Systems, Inc.

Mary Breier
Marketing Communications Specialist
Unitech Systems, Inc.

In general, we have become complacent with accepting errors as a daily occurrence and a cost of doing business. It is accepted that mistakes happen, data errors occur, and systems and processes are not perfect. Accurate, consistent, and reliable information—also known as information integrity—seems to be a luxury that few companies find time to achieve within and around their business.

In the communications industry, the cost of revenue leakage caused by the lack of accurate, consistent, and reliable information is significant enough to impact bottom-line profitability by as much as 15 percent[1]. Network capacity provisioning and activation lack controls; billing systems are unaware of services being used; accumulated call detail record (CDR) errors are increasing; rating tables are applied incorrectly—and this is merely touching the surface of the problem.

This paper explores the concept of information integrity within the context of revenue assurance. It outlines how to approach the problem in a strategic manner. In addition, it describes how it is possible to provide an autonomous, independent, fully auditable, end-to-end control architecture that gives an organization the capability to detect and pinpoint errors at their origin—when they happen, not after the event. How this is accomplished—through an "early warning system" that alerts the business at the first sign of trouble—is also presented.

If you've never had to complain at a supermarket that the price charged was different than the price on the label, if you've never had a transaction appear on your bank statement that wasn't yours, if you've never had a wrong charge show up on your telephone bill, or if you simply trust information to be 100 percent accurate all of the time, then you should stop reading now. Also, if you are one of those rare people who have never received a bill for a service that you've been provided, then keep quiet!

The Current Communications Industry Situation

During the past five years, most communication company strategies have either actively and aggressively captured customers to maximize market penetration and share or, if you were an incumbent provider, protected and minimized customer churn.

New products are launched almost daily on the unsuspecting public, and customers are enticed from one service provider to the next by the promise of better deals. This constantly changing and complex environment sacrifices quality for quantity.

In an attempt to reduce costs, the communications industry has been downsizing. The *Financial Times* estimates that more than 485,000[2] people have been affected by downsizing efforts in the communications industry between April 2001 and January 2002. Communications companies' share prices have plummeted. Outsourcing is the current trend and a way to "pass off" underlying problems. The number of call centers is increasing to handle the barrage of customer complaints. If you believe the estimated 5 percent to 15 percent revenue-leakage figures, then the top communications companies in the United States alone are hemorrhaging from $15 to $47 billion in revenue per year[3]. The morale within the industry is at an all-time low.

During the next five years, the capital, technological, and intellectual investment required within the wireless industry to develop third-generation (3G) networks and products is massive. Only a few organizations will have the power and stamina to survive and create the communications industry of tomorrow. Those few must have strong supportive cultures, solid leadership, effective operational processes, innovative products that penetrate new markets and address their customers needs (quality being one of them).

Revenue Assurance

Revenue assurance should be seen by top executives as a way to provide a competitive advantage for their organization. If approached correctly, revenue assurance becomes a discipline that provides a long-term, repeatable mechanism for capturing revenue that would have been lost otherwise and not realized in the general ledger. Furthermore, revenue assurance can improve customer satisfaction, reduce operational overheads and reduce the risk of damaging the company's brand through bad publicity.

Revenue-Assurance Responsibility

Everyone in an organization should be responsible for revenue assurance. However, from experience, most organizations have many "Revenue Assurance" managers located in several areas—within the network department, others in the information technology (IT) organization, some working specifically in applications such as billing, others in audit, and, in some instances, within finance. Very few organizations have an overall "Head of Revenue Assurance" or a dedicated department that report's directly to the chief executive officer or president of the company.

A department that handles revenue assurance has the potential to increase revenues by 15 percent, making it a critical business function that needs authority and empowerment to instigate changes enterprise-wide. The only way it will succeed is through ownership of the entire organization's revenue-assurance program (see *Figure 1*).

In larger organizations that have a truly global presence, this role often exists both at the group level overseeing the subsidiary parts of the organization and within the subsidiaries.

FIGURE 1

Organization Position for Revenue Assurance

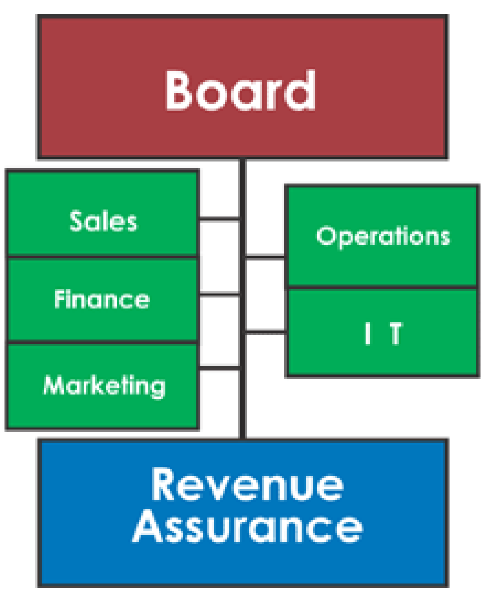

Revenue-Assurance Scope

The scope and impact of revenue assurance can vary depending on business need, but the primary focus to minimize revenue losses and charge inaccuracies and system failures remains consistent.

Revenue assurance goes well beyond simply checking the bills issued and seeing that they are paid. It covers the beginning-to-end process of monitoring every contributing event. We all know the phrase, "garbage in, garbage out." There is no point to having a great billing system if you fail to capture all of the services and products that you want to bill.

Accurate information from the switch to the bill; consistency between applications such as provisioning, billing, and general ledger; and reliable information that can be monitored from the start to the end of the entire process are all within the remit of "revenue assurance."

The End-to-End Operational Process

Within a network, a host of technologies are employed: multiple switches often produced by different manufacturers, network cards that split voice from data to optimize routing, dual-service backbones, and digital subscriber line (DSL)–enabling equipment are but a few. In the near future, these technologies will be expected to provide and manage not only usage data, but also content.

Before a person or company can access a network, he or she goes through a sales process that cumulates in the provisioning and activation of the products and services required from the service provider (see *Figure 2*).

The data generated within the network is the main source of revenue for a communications provider—from the activation and provisioning of new circuits to the generation of CDRs.

CDRs come in a variety of shapes and sizes but generally contain the same information:

- Date and time
- Qualifier
- Calling number
- Called number
- PRI ID number
- B-channel number
- Time of call to disconnect
- Text

A CDR is generated for every call made, even those that don't connect. The lifecycle of an individual CDR is varied and often ends soon after it begins.

In instances where two or more CDRs are generated for one call, suspense processing occurs—with CDRs waiting for the terminating CDR to be matched to the call-starting CDR and any in between. Once the CDRs are created, they need to be put into a standard format. Mediation devices collect the CDRs and reformat them into a standard layout. Some CDRs are assumed to be non-billable, or never billable, and are discarded. Others may be held for processing later, and

FIGURE 2

A Simplified, End-to-End Process Flow

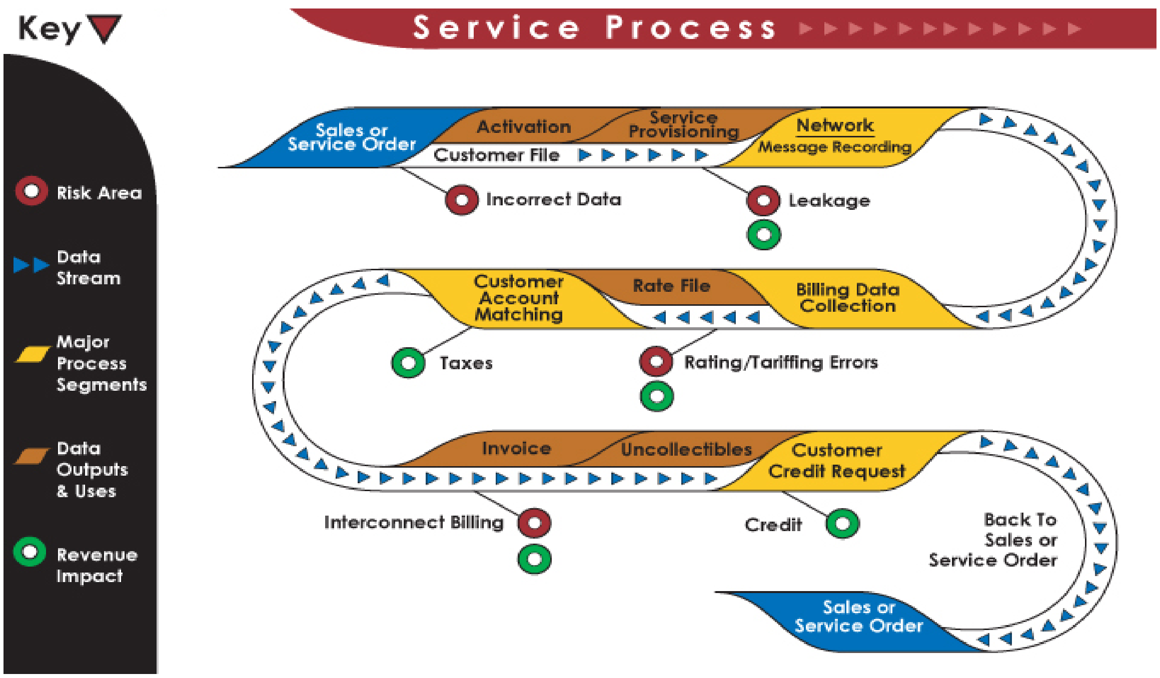

the remaining CDRs are forwarded for rating. Sometimes rating takes place within billing, and, at other times, it is handled as a separate process.

The rating of CDRs is dependent on many factors, such as the time of a call and a customer's call plan. Multiple rating tables exist, often with contradictory or incomplete information. Many service providers allow customers to view rating tables on-line.

To carry out billing, calls are matched to customers, and the information is pulled together into a billing format suitable for delivery as either a printed hard copy or in electronic form to the customer. Therefore, bill verification is critical, since it is the last point at which the service provider can prevent sending an inaccurate bill to the customer.

As soon as the financial information is available from billing, it is forwarded to Finance in the Accounts Receivable revenue stream. Given the complexities, interdependencies, and size of the systems involved, it is not surprising that a change in one area could have an unforeseen damaging impact elsewhere within the process.

Among the areas where revenue can be lost are the following:

- Switch/mediation/billing
- Activation and provisioning
- Rating tables
- Bill invoicing
- Interconnect billing
- Feeds into general ledger

Here are some examples of the types of information errors that have caused lost revenue throughout these areas:

Missing Switch Sequence Numbers: A service provider that was carrying out maintenance on one of its many switches accidentally turned off the sequence capture number. The mediation process was designed to discard CDRs that did not have a sequence number into an error file. This discrepancy wasn't identified until independent, autonomous controls were implemented from the switch on the mediation input file.

Activated Circuits Not in Use: Records that had not been updated resulted in circuits being logged as "in use" when they were actually discontinued and available for reprovisioning and activation. By introducing controls around these processes, the service provider was able to regain capacity it believed was in use. This provided the company with a potential savings of more than $100 million.

Billing Invoices Not Adding Up: A service provider implemented controls to validate whether a bill's contents tallied correctly around their end bill. Previously, customers complained that bill totals were not adding up to the items shown on their bills.

Zero-Rating Applied to CDRs: Toll-free, or free phone, numbers have a zero-rating applied to their numbers. Rating tables were found to have a zero-rating applied to billable numbers.

Incorrectly Formatted CDRs: CDRs generated by a new switch were in a format not recognized by the receiving mediation device and went unrecognized for several months.

New Products Not Recognized: A wireless service provider launched a new call plan only to recall it immediately after its launch due to the inability to identify the product with its supporting applications Without the supporting applications, the new call plan's customers would not be able to use their mobile phones.

Duplicate Records Processed: A new billing system pulled formatted CDR files from the mediation device while the mediation device continued to push formatted files to the billing system. This caused duplicate files that were in turn processed onto the customer's bill. This error didn't go undetected for long!

General Ledger Accounts Receivable Feeds Lost: After billing for revenues, it's important to ensure that journal entries are transferred and processed correctly in the general ledger. For one service provider, these feeds were lost. The ability to accurately report on cash flow is essential for any service provider's operation.

Interconnect Billing Disagreements: A regulator was asked to look into a disagreement between two service providers. The claim was made that 20 million minutes of billable revenue had not been paid by one of them. The regulator sided with the service provider who had the best controlled interconnect process.

These are just a few examples of many situations that could have been avoided with enterprise-wide planning and end-to-end controls.

Strategic Intent

It is possible to achieve an enterprise-wide set of controls and measurements. The approach does not suggest trying to "boil the ocean." Instead, it is advised that a phased approach will help to ensure enterprise-wide, accurate, consistent, and reliable information, which in turn will help to maximize the organization's revenue assurance.

To deliver an effective revenue-assurance strategy, it is necessary to implement a highly visible control and reporting framework supported by a mechanism that is as automated, flexible, and scalable as possible.

This need is driven by the dynamic nature of the communications business, which forces frequent changes to prices and tariff offerings and often results in new or enhanced technical functionality.

The resource constraints and physical restriction of systems often mean that revenue-assurance controls and reporting requirements are scoped into later phases of projects or lost altogether. This results in gaps in end-to-end control and visibility. Revenue can then "leak" from the systems without detection.

In addition, the current fragmented approach of developing individual systems to include control, reporting, and regulatory requirements results in high-cost, multiple developments. This is neither cost- nor resource-efficient.

Revenue-assurance compliance ensures that operational areas apply and maintain appropriate levels of control and produce metrics to set standards. The strategic intent is to automate the end-to-end tracking and reporting mechanism and to reduce system developments across individual components in the end-to-end process (see *Figure 3*).

Performance measurements should be analyzed for trends and key metrics, and then produced for the board to review. A fully automated system would provide results via a Web-based browser.

Where possible, revenue-assurance control should be maintained through non-intrusive automated monitoring across all components of the end-to-end core business process.

By implementing an automated autonomous approach, end-to-end monitoring and measurement can take place in an independent "live" environment, where it is possible to delivers robust control and reporting capabilities to the desktop. This eradicates the bottleneck for development on each of the end-to-end systems and facilitates accurate, consistent, and reliable management reporting.

Reaping the Benefits

This approach delivers real benefits to the service provider, including the following:

- Vastly improved control and visibility, and early warnings of error, loss, and leakage for all systems
- The ability to action and recover rejected or suspended call data in a timelier manner and, therefore, stem losses and unnecessary write-offs
- The ability to identify potential losses that are not currently visible at all
- Maintenance of revenue-assurance developments that is performed independently of existing systems
- No bottlenecks or priority conflicts on existing systems
- Less people required for extracting data, performing analysis, and producing reports
- "Live" reports that are available to authorized users at any time via a Web site that gives instant visibility to end-to-end core business performance
- A scalable approach that can be flexibly "plugged" into other systems

Although it is not essential, to reap maximum benefits from this approach, the "Revenue Assurance" function should have enterprise-wide responsibility for monitoring and measuring. In addition, it should be recognized at the highest levels within the organization. This would be the optimal situation.

FIGURE 3

Autonomous End-to-End Controls

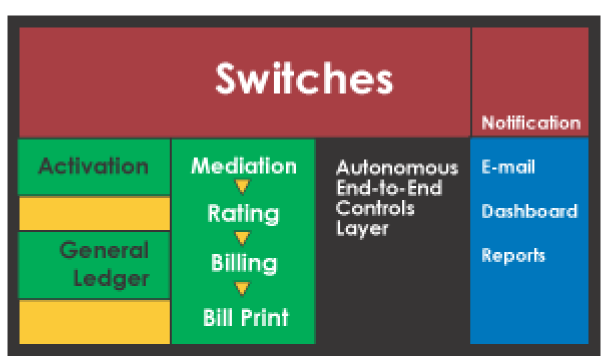

Everyone recognizes the importance and relative costs of a lack of information integrity, but we seem to be too paralyzed by the potential magnitude and scale of the issue to do anything enterprise-wide.

Phase One: Assessment

In assessing risks to information integrity, the objective is to identify and prioritize the information-error exposures in the business process. To do this, the assessor needs to do the following:

- Develop a process flow diagram
- List all potential sources of information errors and prioritize the list based on risk to the business
- Provide suggestions for next steps, including a recommended approach

The methodology for an integrity risk assessment consists of five phases, which are outlined in *Figure 4.* After the integrity risk assessment is performed, a document is delivered that captures the knowledge of employees and details the processes and procedures between various systems and applications (in this instance, it would be focused on the issue of revenue assurance). The knowledge would be documented to provide one complete end-to-end view of the system.

Furthermore, the assessment needs to assess the effectiveness of these procedures by component and where data leakage could be, or is, occurring. These areas of leakage indicate exposures to information errors and risks to information integrity.

An assessment that looks at the end-to-end revenue-assurance system documents and builds an understanding of the full "CDR life cycle" from point of generation on the switch to production of the billing file to another service provider or end customer, and ultimately into the general ledger.

Finally, where possible, the assessment identifies what happens if there is a failure.

Phase Two: Detection

After identifying the potential information error exposures in the business system and process, the next step is to put a series of sensors in place that will detect errors as they occur.

Because there are many different types of information errors in a process, there are a variety of information-error detection sensors that can be designed around the characteristics of each type of error. A combination of each of these sensors is usually required to identify all of the information errors in a given business process.

These error types are grouped into four categories:

- Data Validation
- Transaction Tracking
- Balancing and Reconciliation
- File Validation

Building an autonomous, automated sensor array around the company's processes creates an "early warning system" matrix of controls to detect information errors before senior executives, business partners, and customers see them—and before the errors become front-page news.

Table 1 provides a detailed breakdown of information error types that could affect a business.

Phase 3: Resolution

Sensors that non-invasively monitor the end-to-end process create the controls architecture layer that provides service providers with confidence that their core business processes are operating effectively and efficiently, because any issues are identified at the source, when they occur.

TABLE 1

Error Group Sensor Types

Data Validation	Transaction Tracking	Balancing & Reconciliation	File Validation
Format Content Context Error Rate Profile	Value Time Path	Run-to-Run Point-to-Point Over Time Across Platforms	Usage Sequence Specs

FIGURE 4

Assessment Methodology

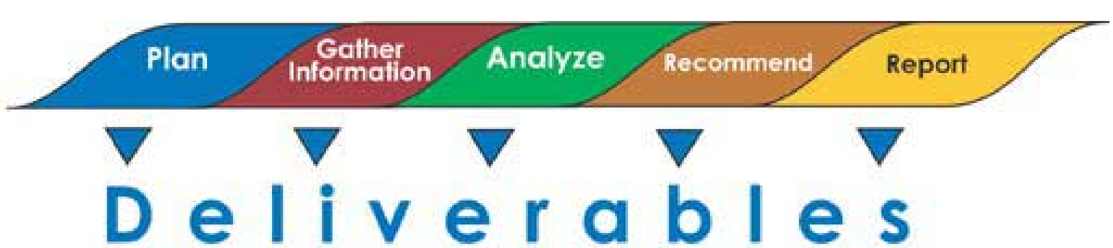

Beyond the detection phase, an integrated set of resolution tools is often desired to round out the functionality of the "early warning system."

Firstly, a dashboard view provides the ability to view errors flagged by the sensors and the current revenue flow status of the entire enterprise-wide environment via a normal Web browser. With the dashboard, service providers are able to drill down into "problem areas" and, with complementary communication applications, automatically e-mail or page the appropriate people about the problem, enabling quick resolution.

Secondly, the resolution phase should allow for the automated correction of errors and the ability to provide a workflow and correction environment for any detected information errors. In essence, the resolution piece of an "early warning system" should provide the opportunity to fix the error, thereby assuring the accuracy, consistency, and reliability of the information.

Phase 4: Prevention

After completing the initial deployment of the detection and resolution phases, focus shifts to ongoing deployment to maximize the value received from the "early warning system." This is the prevention phase.

To determine what additional deployment will provide the greatest return, the prevention phase emphasizes measurement of current value to determine next steps.

The optimization of value from the "early warning system" approach relies on two elements:

1. The measurement of value received from the current deployment
2. A prediction of future value from either enhancing the implementation of tools already in place or through implementing new tools in an upstream or downstream process

Using a series of best-practice valuation and analysis methods, the recommended course of continuous improvement can be determined.

Strategic Initiative Advantages

By effectively introducing an "early warning system" strategic initiative that looks at revenue assurance end to end, a number of comparative advantages are realized. These advantages include the following intangible and conditional benefits:

- Improved decision making by providing timely, comprehensive, and accurate information. For example, bad debts that perhaps "disappeared" from a system are identified, allowing for the recovery of the monies owed.
- Replacement of low-value staff activities with high-value work. For example, a lawyer could end up spending less time building, checking, and presenting price lists/rating tables and more time dealing with "legal" issues.
- Established and centralized responsibility for revenue assurance, which reduces the need for every system to track its own part of this responsibility, improving efficiency and reducing costs.
- Protection of the brand from bad publicity, regulators, and customer complaints. Not only will customer retention increase, but new customer sign-ups may increase also.
- Standardized approach to revenue assurance that can reduce costs.
- Proactive approach toward the identification of errors versus a nonexistent approach.
- Assurance that systems have the ability to grow with the company as it grows.
- Utilized information as an asset. For example, marketing could potentially examine and analyze call usage related to new products and identify trends in calls that suggest new products.
- Assurance that *all* revenues earned are being received—an approach that would be well received by shareholders and, ultimately, the stock market.
- True end-to-end audit ability versus application and system disparate auditing.
- Competitive advantage within a highly competitive environment.

In addition to this tangle of intangible and conditional types of benefits, there are a number of "tangible" benefits that have been generated within communications companies[4], including the following:

- A 7 percent increase in billable transactions
- A 95 percent reduction in double-posting errors
- A 53 percent reduction in error-related costs
- A reduction from two weeks to only one day to perform a monthly financial close
- An increase in available network capacity

Once again, effectively addressing critical issues such as revenue assurance requires a strategic approach, good planning, and decision-making authority.

The Information Integrity Coalition

The information age to date has been about the technology that handles the data. The communications industry has played a key role in its ascent. The Information Integrity Coalition, formed in 2001, is looking into ways to address those areas beyond the technology.

According to Madhavan K. Nayar, one of the founders of the Information Integrity Coalition, the time has come to make the information age about information[5]. In outlining the economic impact of information errors, Mr. Nayar acknowledges, "We are aware errors occur; we are resigned to them as a cost of doing business; and we are skeptical that we can do anything about them." However, using current revenue-leakage indicators, he underscores that resignation is not the answer.

The Information Integrity Coalition takes a holistic approach to a solution, seeking nothing less than the formation standards and a new industry. That new industry will be one that Mr. Nayar says will address information integrity as more than a simply technical issue—as something multidimensional. He outlines several inquiry possibilities: data (numbers, graphics), assurance methods (security, audit, controls), assurance activities (prevention, detection, monitoring, correction), system life phases (design, develop, implement), system types (manual, online, interactive), industry-specific (banking, health, education, insurance, communications), business (cost, risk, opportunity), and disciplines (statistics, economics, psychology, ethics).

There are companies out there right now that address some of these information-integrity dimensions, such as security, error detection, information monitoring, industry-specific solutions, and data companies.

Their efforts, however, are fragmented. The Information Integrity Coalition seeks to reconfigure them through the creation of an information-integrity discipline that is complete with standards.

Summary

Communications service providers are seeking ways to become more competitive, efficient, and effective in their daily operations. Revenue assurance is needed within the communications industry, because the cost of revenue leakage caused by the lack of accurate, consistent, and reliable information is significant. Through a holistic, strategic, and phased approach, service providers can introduce a standard way to monitor, measure, and control their core business processes related to revenue assurance. Improved integrity of information is only one result of implementing this approach.

Through this consultative, strategic approach, partly consisting of four phases—assessment, detection, resolution, and prevention—service providers can implement an "early warning system" for information errors so that they can be corrected before they affect customers, the business, or become headline news. This approach helps to create information integrity in the enterprise and helps in the effort to assure that all revenue is recognized.

Few organizations have a true enterprise-wide view of revenue leakage. This is not good business, and, while the task ahead is far from easy, you must (to quote Churchill) "never, never, never, never, never, never, never, never, never…give up!" A company's information is its biggest undervalued and unrealized asset. A board-level–endorsed revenue-assurance strategy is the first step in the right direction.

Notes

1. Michelle Teofan, *Controlling the Wireless Revenue Web, White Paper World Class Billing Practice,* KPMG.
2. http://news.ft.com/ft/gx.cgi/ftc?pagename=View&c=Article&cid=FT3MOCS3OPC&live=true&tagid=ZZZPCGI2B0C&subheading=telecoms [accessed January 15, 2002].
3. http://specials.ft.com/ft500/may2001/FT396T1OKMC.html (Sector Codes 673 & 678)
4. Taken from confidential company applications for the Excellence in Information Integrity Award.
5. Michele Fitzpatrick, *Information Integrity Coalition Launch*, October 22, 2001.

Backhaul from International Cable Systems: A New Carrier Market Opportunity

John Gudgel
Senior Marketing Manager, Global Carriers
Lightscape Networks

I. Introduction

An optical tidal wave is hitting the shores of the United States. By the end of 2002, at least 24 synchronous digital hierarchy (SDH) international cable systems will have landed in the United States and become operational (see *Figure 1*). The total capacity of these systems is more than 20 terabits per second (Tbps)—a bandwidth sufficient, in theory, to allow every man, woman, and child in the United States to simultaneously make an overseas phone call.

In the trans-Atlantic corridor alone, capacity in the next year will increase by more than 500 percent. This growth is spurred by the demands of hundreds of new global carriers and Internet service providers (ISPs) hungry for international bandwidth to satisfy their Internet backbone demands.

However, international bandwidth is only useful if it can be delivered to the end user. Once an international cable is landed at a cable station, it still has to be extended from the station to a carrier's international point-of-presence (POP) and, from there, via domestic networks, to the end user. This cable system backhaul can be the Achilles heel of international capacity delivery, significantly increasing both the cost and provisioning time to implement service. Yet for those carriers that understand the technical and economic challenges associated with backhaul and that proactively implement backhaul solutions, there is a potentially large market opportunity.

This paper describes the current international cable environment, identifies the technical and economic issues associated with cable-system backhaul in the United States, provides some possible backhaul solutions, and analyzes the potential market opportunity associated with the lease or sale of backhaul facilities.

II. Submarine Cable Systems—A Brief Tutorial

The first international submarine cable system connecting England and France was laid in 1850. Used for telegraphy, it consisted of a single insulated copper wire that could transmit only one message at a time. Unfortunately, the cable was cut by a fisherman less than two days after going operational.

Undersea cables have come a long way since 1850. Most of the undersea cables being installed today are fiber-optic ring-protected systems capable of handling hundreds of thousands or even millions of simultaneous voice calls. The technology that makes this possible is dense wavelength division multiplexing (DWDM) (see *Figure 2*). This technology allows multiple high-speed channels (currently 2.5 of 10 gigabits per second [Gbps]) using different wavelengths (currently as many as 40 colors) of light to be transmitted over a single fiber. Most international submarine cables being installed today consist of two, four, or eight fiber strands, typically configured in a ring configuration. The formula to calculate the carrying capacity of a cable system is as follows:

- # strands of fiber x # channels of DWDM x size of the DWDM channel x the type
- of configuration (one for point-to-point, two for a ring configuration)

Thus the carrying capacity of TAT-14 is as follows:

- 8 strands of fiber x 32 channels per fiber x 2.5 Gbps per channel x 2 (ring) = 1,280 Gbps
- or 640 Gbps (fully protected)

All international cables are currently being implemented using International Telecommunication Union (ITU) standards that define optical transmission using the SDH. This

FIGURE 1

International Submarine Cables Landing in the United States

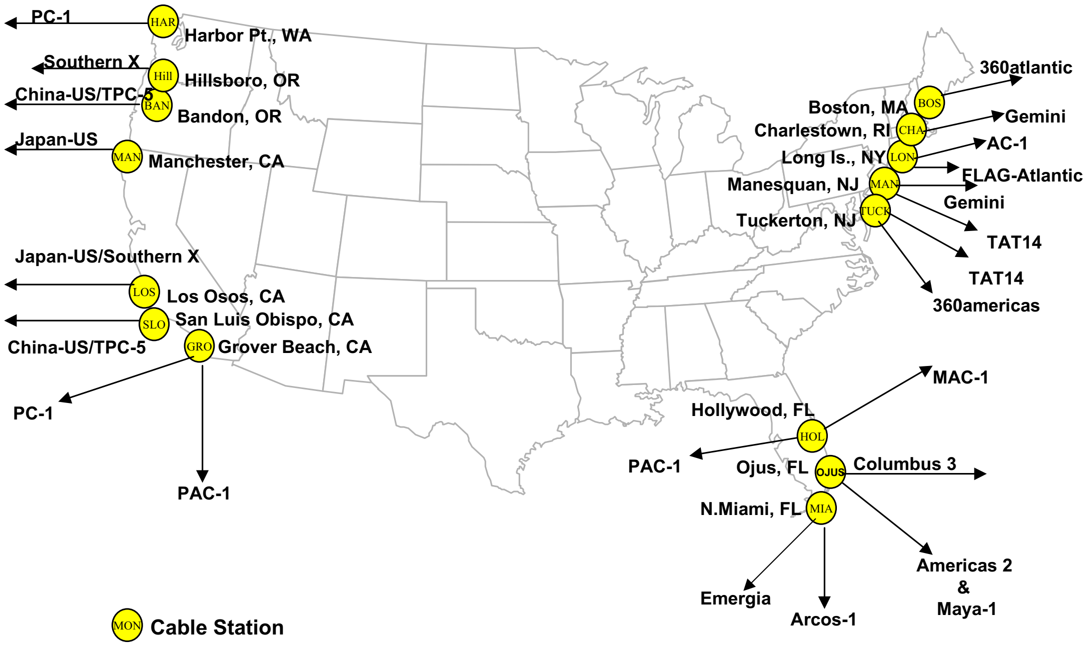

optic transmission standard is different from the synchronous optical network (SONET) standard, which is used by domestic carriers in the United States. This difference is an important technical consideration when implementing international circuits, because oftentimes this capacity must be converted from SDH to SONET to be transmitted over networks in the United States.

The basic unit of SDH capacity typically utilized in international cables systems is the synchronous transfer mode (STM)–1, which is a 155 megabits per second (Mbps) circuit. Its SONET equivalent is an optical carrier (OC)–3. Oftentimes, cable systems are sized based upon their number of STM–1 equivalents. Thus TAT-14, which has a carrying capacity of 1,280 Gbps (640 Gbps ring protected), can support as many as 8,192 unprotected or 4,096 protected STM–1s.

Until a few years ago, virtually all international cables were built by consortiums consisting of a large number of global carriers that agreed to share the high cost of building the system. The carriers sign a construction and maintenance agreement (CMA), which defines the amount of capacity that each carrier owns in the system and the terms and conditions for participating in the cable construction and ongoing operation. Cable system capacity is purchased as an indefeasible right of usage (IRU). An IRU is like owning a condominium in a building. The owner pays an upfront capital price for the right to use a specific amount of capacity (e.g., an STM–1) for the lifetime of the cable (typically 15 or 20 years) and ongoing annual maintenance fees (typically 5% to 7% of the initial cost). Under the traditional consortium construction model, the IRU is only for the submarine portion of the capacity (cable station–to–cable station). The terrestrial portion, or backhaul, which connects the capacity from the cable landing station back into the owner's domestic network, is not part of the IRU. This backhaul must be leased or purchased separately at considerable effort and expense.

Recently, a number of new cable systems have been financed and built by ***joint ventures*** (e.g., Gemini, which is a joint venture of WorldCom and Cable & Wireless) or by ***private cable companies,*** such as Global Crossing. These new construction entities are changing the rules by which capacity is purchased. Among the most important changes is that IRUs are now being sold ***city-to-city,*** with a large portion of the backhaul at both ends being included in the IRU. However, for the time being, a large number of cable systems continue to be built by consortiums, cable station–to–cable station, and still require that the capacity owner make their own arrangements for backhaul.

III. Cable System Overview

Table 1 is a list of the U.S.–terminating SDH cable systems currently in operation or under construction. By the end of 2001, these new SDH cable systems will increase the total capacity of bandwidth coming into the United States by more than 5,000 times versus pre-1995 levels. Subsequently, the price of IRU capacity has plummeted from more than $20 million per cable station–to–cable station trans-Atlantic STM–1 in 1995 to, in some cases, less than $500,000 for a similar city-to-city STM–1 in 2001.

Trans-Atlantic Cable Systems

At the end of 2000, there were five SDH cable systems operational between the United States and Europe. They are TAT-12/13 (which went operational in 1995), Gemini and AC-1 (which came into service in 1998), and Columbus 3 and AC-2 (which came into service in late-2000). TAT-12/13 is a ring system that connects the United States with France

FIGURE 2

DWDM allows one to split and combine different wavelengths in a single fiber.

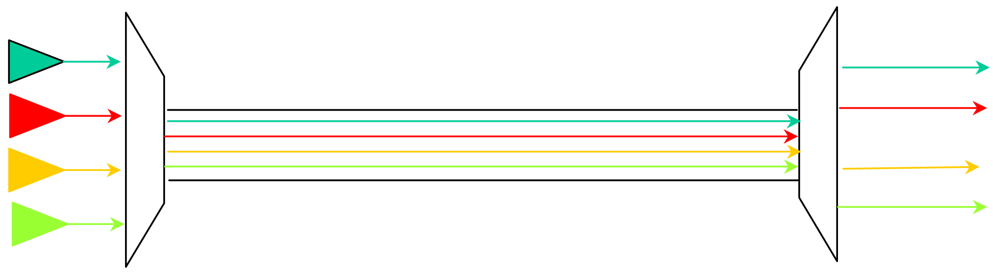

TABLE 1

U.S.–Terminating SDH International Cables

Cable System	Venture Type	Region	Current Est. RFS Date	Type	System Capacity in Gbps	System Capacity in STM–1s	City-to-City?	Web Site
360americas (former Atlantica)	Private	Atlantic/ Caribbean	2001	SDH Ring	1,280	8,192	Yes	*www.360.net*
360altlantic (former Hibernia)	Private	Atlantic	2001	SDH Ring	1,280	8,192	Yes	*www.360.net*
360pacific	Private	Pacific	2002	SDH Ring	4,800	30,720	Yes	*www.360.net*
Americas 2	Consortium	Caribbean	In Service	Festoon	80	512	No	N/A
Apollo	Private	Atlantic	2002	SDH Ring	3,200	20,480	Yes	N/A
Arcos 1	Private	Caribbean	2001	SDH Ring	960	6,144	No	*www.arcos1.com*
Atlantic Crossing (AC-1)	Private	Atlantic	In Service	SDH Ring	80	512	Yes	*www.global crossing.com*
Atlantic Crossing 2 (AC-2)/ Level 3 Yellow	Joint	Atlantic	In Service	SDH Ring	1,280	8,192	Yes	*www.global crossing.com*
China-US	Consortium	Pacific	In Service	SDH Ring	160	1,024	No	N/A
Columbus 3	Consortium	Atlantic/ Caribbean	In Service	SDH Ring	40	256	No	N/A
Emergia	Private	Caribbean	2001	SDH Ring	1,920	12,288	Yes	*www.e-mergia.com*
FLAG-Atlantic	Joint	Atlantic	2001	SDH Ring	2,560	16,384	Yes	*www.flag atlantic.com*
FLAG-Pacific	Joint	Pacific	2002	SDH Ring	5,120	32,768	Yes	*www.flag pacific.com*
Gemini	Joint	Atlantic	In Service	SDH Ring	30	192	Yes	*www.gemini.bm*
Japan-US	Consortium	Pacific	2001	SDH Ring	640	4,096	No	*www.japan-us.org*
Maya-1	Consortium	Caribbean	In Service	Festoon	20	128	No	*www.maya-1.com*
Mid-Atlantic Crossing (MAC-1)	Private	Atlantic/ Caribbean	In Service	SDH Ring	20	128	Yes	*www.global crossing.com*
Pacific Crossing (PC-1)	Private	Pacific	In Service	SDH Ring	80	512	Yes	*www.global crossing.com*
Pan American Crossing (PAC-1)	Private	Caribbean	In Service	Festoon	40	256	Yes	*www.global crossing.com*
South American Crossing (SAC-1)	Private	Caribbean	2001	SDH Ring	80	512	Yes	*www.global crossing.com*
Southern Cross	Joint	Pacific	In Service	SDH Ring	160	1,024	No	*www.southern crosscables.com*
TAT-12/13	Consortium	Atlantic	In Service	SDH Ring	10	64	No	N/A
TAT-14	Consortium	Atlantic	2001	SDH Ring	640	4,096	No	*www.tat-14.com*
TPC-5	Consortium	Pacific	In Service	SDH Ring	10	64	No	N/A

and the United Kingdom. It recently underwent an in-system upgrade that increased its ring-protected capacity from 5 to 10 Gbps. Gemini and AC-1 are larger ringed systems that broke the virtual monopoly that TAT-12/13 had on trans-Atlantic capacity from 1995 to 1997. Gemini and AC-1 have been competing with each other since 1998. In the process, they drove the price down for a trans-Atlantic STM–1 to less than $4 million per IRU. In late-2000, two new systems—AC-2 and Columbus 3—were completed, further intensifying trans-Atlantic bandwidth competition. However, all of these systems will be facing much stiffer competition when several new terabit-capacity trans-Atlantic cables are completed in mid-2001 and 2002.

The new trans-Atlantic systems include TAT-14, FLAG Atlantic, 360americas (formerly Atlantica), 360atlantic (formerly Hibernia), and Apollo. When all completed, the amount of trans-Atlantic ring-protected capacity will have increased by more than 50,000 STM–1s. This may produce a glut of capacity and a buyers market as cable-system owners scramble to sell IRUs before the price drops out of the market. Already, prices have fallen to unprecedented lows. However, many companies are gambling that the demand fueled by the Internet will quickly consume the excess supply. As a result, existing cable system owners are contemplating in system upgrades, and even larger multi-terabit systems are being discussed.

Trans-Pacific Cable Systems

For the past several years, the capacity situation has been very bleak. Between 1995 and 1999, the only SDH cable system with any IRU capacity was TPC-5, and that bandwidth, when you could find it, was very expensive. An in-system upgrade of TPC-5 from 5 to 10 Gbps (ring protected) in 1999 only relieved the trans-Pacific capacity famine slightly. The situation improved dramatically in early 2000, when two systems China-US and Pacific Crossing 1 (PC-1) went into operation. In addition to the cables that have been described, several other SDH cables are schedule to go into full operation in 2001. These systems include Southern Cross, which will connect the United States with Australia and New Zealand, and Japan-US. In 2002, two new terabit systems, 360pacific and FLAG-Pacific are also scheduled to go into service. Further information on all of the trans-Pacific cable systems in given in Table 1.

Cable Systems Connecting the United States to South and Central America

There are a surprising number of new cable systems being implemented that connect the United States to Central and South America and to the islands in the Caribbean. These systems include Americas 2, Maya-1, Arcos, Emergia, and South American Crossing 1 (SAC-1). All of these systems land in the United States in Southern Florida. In the case of SAC-1, this landing is actually via another Global Crossing system: Mid-Atlantic Crossing (MAC), which is a feeder cable that runs from Florida to Bermuda (the actual SAC-1 landing site) and then to the AC-1 landing site on Long Island, New York. Another Global Crossing cable system, Pan-America Crossing -1 (PAC), will eventually connect from the Florida landing site for MAC-1 to the California cable station for PC-1. At that point, Global Crossing will be able to provide transit service between all of its U.S.–terminating cables: AC-1, MAC-1, SAC-1, PAC-1, and PC-1.

IV. Backhaul Issues and Alternatives

When an international submarine cable is landed in the United States, it is natural to assume that it is a fairly easy matter to extend that capacity on-shore from the cable station into a carrier's domestic network and, from there, to the end users' location. After all, fiber-optic transmission is in wide use in the United States, the circuits run over dry land versus under the sea, and one would think that any carrier with fiber capacity and a little bit of ingenuity could easily handle the task. In fact, cable-system backhaul implementation can, at times, be far from easy. There are several technical and economic issues that can make backhaul difficult and expensive to implement.

To begin with, all international cables currently are being implemented using SDH technology. This means that capacity on these systems is being transmitted as SDH, lower-rate circuits must be mapped onto higher-rate SDH circuits, and the hand-off of the circuits at the U.S. cable station is also SDH. Because the standard for fiber-optic transmission in the United States is SONET, delivering this capacity to the end user can be a challenge. More and more, end-user equipment, such as routers and asynchronous transfer mode (ATM) switches, are being designed to support both SONET and SDH interfaces. The problem is getting the SDH circuit from the cable station to the equipment. One obvious solution is to convert the SDH capacity to SONET using some type of gateway protocol converter, multiplexer, or digital access cross-connect system (DACS). These gateway boxes can perform the SDH to SONET conversion (e.g., STM–1 to OC–3) and, in some cases, can also map U.S. standard plesiochronous digital hierarchy (PDH) capacity (T1, digital signal [DS]–3, etc.) onto international SDH circuits.

Another problem with backhauling international circuits is economic. Up until a few years ago, nearly all international submarine cables landing in the United States were terminated at stations belonging to AT&T. AT&T controlled access to these stations and, as a result, had a virtual monopoly on the international cable backhaul business in the United States. Carriers owning capacity on U.S.–terminating cables had to deal with AT&T, which set the rules and costs for activating backhaul capacity. The rules, which were bureaucratically adhered to, could cause lengthy provisioning delays oftentimes resulting in backhaul intervals of six months or more. Furthermore, AT&T charged considerably more for SDH backhaul capacity versus comparable domestic SONET circuits.

Recently, the AT&T backhaul monopoly has been broken. Systems such as Gemini and AC-1, which came into service in 1998, sell city-to-city service. City-to-city service includes the backhaul from cable stations to select major metropolitan area "city sites," which are usually major carrier hotels such as 60 Hudson and 111 8^{th} Avenue in New York City. Examples of other city-to-city systems being implemented include FLAG-Atlantic, MAC-1, PAC-1, PC-1, and 360atlantic. The problem with city-to-city service is that it is only available to certain buildings in specific cities. If users need to extend capacity to a building or city that is not covered, they may still have a problem. They still need to extend an SDH circuit from the city site to their location. In some instances, certain competitive access providers (CAPs) and interexchange carriers (IXCs) are offering SDH service to select locations. Thus, for example, some carriers may be offering leased STM–1s from a city site to a major Internet network access point (NAP), such as MAE East. Or, the operator of the city site may offer SDH–to–SONET conversion that will allow, for instance, the delivery of an STM–1 as an OC–3. It is also possible that a city-site operator will allow customers to co-locate a gateway box so they can do the SDH–to–SONET conversion themselves.

Many international cables continue to be implemented cable station–to–cable station, without any in-system provisions for backhaul. However, most cables now being built allow owners the right to co-locate at the cable station and to provide their own backhaul via dark fiber. For example, several carriers, including Sprint and Williams Communications, have co-located at the Bandon, Oregon cable station for China-US. They have implemented backhaul for their own needs and, also, are competing with AT&T by leasing backhaul facilities to other carriers. This limited competition has driven the monthly recurring cost for a Bandon-to-San Francisco eased backhaul STM–1 from more than $100,000 per month to less than $40,000 per month. Other cable station–to–cable station cables allowing co-location include TAT-14, Japan-US, Southern Cross, Americas 2, and Maya-1.

Many carriers, especially those owning large amounts of capacity on these cables, are actively arranging cable-station co-location and the implementation of their own dark fiber backhaul. They are implementing this service not only to save money by not having to pay another carrier for leased backhaul, but also to speed provisioning time and generate additional revenue by being able to sell backhaul themselves. To implement this service, they are installing SDH concentration and DWDM transmission equipment at the cable station and connecting it via dark fiber to similar equipment located in one or more international access POPs. At the access POP, they may have a gateway box to do SDH–to–SONET conversion so that capacity can be delivered to customers. They may also have international voice or ATM switches, routers for Internet access, or co-location space for customers wanting direct access to international facilities. Such a configuration is shown in *Figure 3*.

Thus, carriers that plan to purchase international capacity have a number of alternatives to consider regarding how to backhaul that capacity in the United States. They can purchase a city-to-city IRU and then only have to worry about the last mile between the city site and their POP. They can buy traditional cable station–to–cable station capacity and then either lease backhaul capacity from another carrier or build their own dark fiber backhaul facilities. In any case, they may still need to convert that capacity from SDH to SONET or break the SDH capacity down into DS–3s or T1s so it can be delivered to an end user. They can do this by either purchasing a gateway box to do the SDH/SONET conversion and mux/demuxing or pay another carrier for these services.

V. "Backhaul" Market Opportunities

With the tidal wave of new cables landing ashore in the United States, the proliferation of fiber-optic cable around the country, the willingness of cable-station operators to allow co-location, and the economics of DWDM, a number of opportunities have arisen for carriers to sell backhaul-related products.

FIGURE 3

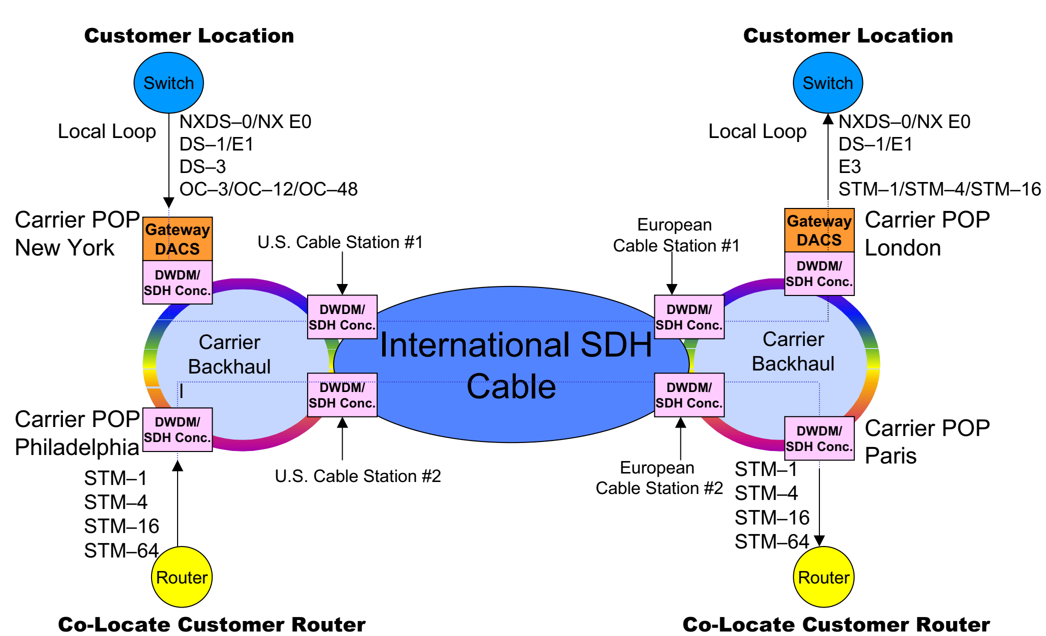

The first obvious product is the lease or sale of backhaul facilities. Carriers that build their own backhaul networks can lease STM–1s, STM–4s, STM–16s, or whole wavelengths to other carriers. They can also sell these same facilities to carriers as backhaul IRUs at a fixed capital price plus annual recurring maintenance.

Another potential product is transit between two cable systems. This product includes the lease or sale of backhaul from two cable systems (e.g., TAT-14 and Japan-US), as well as a long-haul circuit connecting the two backhaul segments. Once again, these facilities could be leased to a carrier on a month-to-month or term basis, or sold to the carrier as a transit IRU.

Carriers that own backhaul networks to multiple cable systems can also provide restoration services to carriers or to cable-system operators. Under this scenario, backhaul network operators would reserve capacity for customers to be used in the event that a cable system suffers a catastrophic failure that disrupts service. The reserve backhaul capacity would be used to switch a customer's capacity from one cable serving a region (e.g., a Gemini STM–1 in the Atlantic) to another cable serving the same region (e.g., AC-1). In today's age of self-healing SDH cables, many carriers may see this service as unnecessary. However, several cables currently are being deployed initially as collapsed rings or point-to-point systems without restoration, and failures have occurred on supposedly "fully protected" submarine cables. Restoration services can be sold either in advance as a kind of insurance policy against possible failure or "by the drink" during an emergency situation.

Finally, carriers can sell backhaul bundled with other products and services. For example, a carrier that owns facilities at both ends of a cable can offer one-stop-shop, city-to-city leased international private lines or IRUs that can be quickly provisioned for customer use. Delivery of these circuits might include optional bundled SDH/SONET conversion, mux/demux of capacity into lower-rate domestic circuits, or a high-speed Internet port. Thus, backhaul can be used to enhance a carrier's entire portfolio of global products.

VI. Conclusions

The coming year will offer global carriers unprecedented access to trans-oceanic cable capacity. However, those carriers will need to plan ahead to ensure that they have a backhaul solution in place so they can access their international capacity when it becomes available.

These carriers need to ask themselves what they plan to use the capacity for and where it will need to be delivered. Once these questions are answered, only then can a carrier adequately determine the best way in which to backhaul this capacity. Their backhaul solution could be a lit solution that is leased or purchased from another carrier, or one that involves cable-station co-location and dark fiber. They may also need to determine whether or not they require SDH–to–SONET conversion or the breakdown of their capacity into smaller increments. If they do, they need to install a gateway DACS or some other gateway solution.

Lastly, once a carrier has a backhaul solution in place, they can then pursue a number of potential business opportunities that can not only produce direct revenue, but also enhance their entire global product portfolio.

This paper was originally published in the July 2001 issue of Lightwave Magazine.

Simulation of the Harmful Consequences of Self-Similar Network Traffic: A Model for Network Engineers

Tibor Gyires
Professor, Applied Computer Science
Illinois State University

Abstract

For a long time, the area of network traffic research has lacked adequate traffic measurements. During the past years, a large amount of network traffic measurements have become available, collected in the Web and high-speed networks. Measurements of real traffic indicate that burstiness is present on a wide range of time scales and show evidence of self-similarity and long-range dependency. Many traditional models originating in conventional voice networks either assumed that the underlying time series are short-range dependent or do not focus on the various, interrelated network components affected by bursty traffic. To provide explanations for empirically observed phenomena, such as burstiness in terms of physical network parameters, we need to develop models that can allow for the details of the complex architecture of today's networks and can analyze the interrelated network parameters that ultimately determine the performance and operation of a network. In the paper, we investigate the harmful consequences of burstiness on various network components using a model implemented by a discrete event simulation methodology. The methodology based on a modified version of the M/Pareto model can measure network performance parameters, such as the momentarily utilization of links, response time, and the queuing performance of interconnected switches and routers along the traffic paths. Our paper intends to narrow the gap between existing, well-known theoretical results and their applicability in everyday, practical network analysis and modeling. Our methodology can help network designers and engineers, the ultimate users of traffic modeling, to understand the dynamic nature of network traffic and assist them to design, monitor, and control complex, high-speed networks in everyday practice.

1. Introduction

Recent measurements of local-area network (LAN) traffic [1] and wide-area network (WAN) traffic [2] have proved that the widely used Markovian process models cannot be applied for today's network traffic. If the traffic were Markovian process, the traffic's burst length would be smoothed by averaging over a long time scale contradicting with the observations of today's traffic characteristics. Measurements of real traffic also prove [3] that traffic burstiness is present on a wide range of time scales. Traffic that is bursty on many or all time scales can be characterized statistically using the concept of self-similarity. Self-similarity is often associated with objects in fractal geometry, objects that appear to look alike regardless of the scale at which they are viewed. In case of stochastic processes such as time series, the term self-similarity refers to the process' distribution-when viewed at varying time scales, the process' distribution remains the same. Self-similar time series has noticeable bursts, long periods with extremely high values on all time scales. Characteristics of network traffic, such as packets per second, bytes per second, or length of frames can be considered as stochastic time series. Therefore, measuring traffic burstiness is the same as characterizing the self-similarity of the corresponding time series.

The self-similarity of network traffic has also been observed in studies in numerous papers, such as [4], [5], and [6]. These and other papers show that packet loss, buffer utilization, and response time are totally different when simulations use either real traffic data or synthetic data that include self-similarity.

1.1. Background

For more details of self-similarity in time series and the associated statistical tests, see [3], [7], and [8]. Our paper follows the definitions of these papers.

Let $X = (X_t: t = 0, 1, 2, ...)$ be a covariance stationary stochastic process. Such a process has a constant mean $\mu = E[X_t]$, finite variance $\sigma^2 = E[(X_t - \mu)^2]$, and an autocorrelation function $r(k) = E[(X_t - \mu)(X_{t+k} - \mu)]/E[(X_t - \mu)^2]$ $(k = 0, 1, 2, ...)$ that depends only on k. It is assumed that X has an autocorrelation function of the form:

$$(1) \quad r(k) \sim \alpha k^{-\beta}, \qquad k \to \infty$$

where $0 < \beta < 1$ and α is a positive constant. Let $X^{(m)} = (X^{(m)}_{(k)} : k = 1, 2, 3, ..., m = 1,2,3,...)$ represent a new time series obtained by averaging the original series X over non-overlapping blocks of size m. For each m = 1, 2, 3,…, $X^{(m)}$ is specified by $X^{(m)}_k = 1/m(X_{km-m+1} + \ldots + X_{km}), (k \geq 1)$. Let $r^{(m)}$ denote the autocorrelation function of the aggregated time series $X^{(m)}$.

1.1.1. Definition of Self-Similarity

The process X is called exactly self-similar with self-similarity parameter $H = 1 - \beta/2$ if the corresponding aggregated processes $X^{(m)}$ have the same correlation structure as X, i.e., $r^{(m)}(k) = r(k)$, for all m = l, 2,... (k = l, 2, 3,...).

A covariance stationary process X is called asymptotically self-similar with self-similarity parameter

$H = 1 - \beta/2$ if for all k large enough $r^{(m)}(k) \to r(k)$ as $m \to \infty$, $0.5 \leq H \leq 1$.

1.1.2. Definition of Long-Range Dependency

A stationary process is called long-range dependent if the sum of the autocorrelation values approaches infinity: $\Sigma_k\, r(k) \to \infty$. Otherwise, it is called short-range dependent.

FIGURE 1

The Self-Similar Nature of Internet Network Traffic

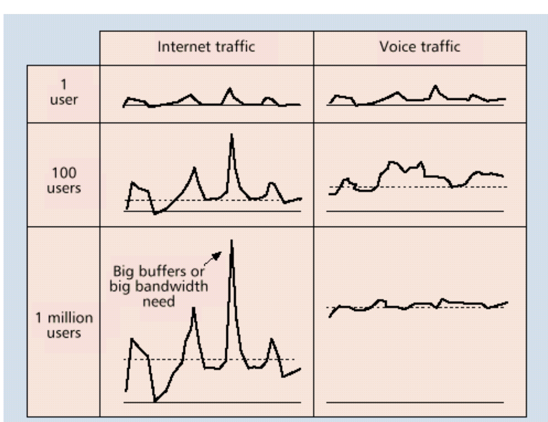

(Source: M. Listanti and V. Eramo, "Architectural and Technological Issues for Future Optical Internet Networks," IEEE Communications (September 2000), 82-92.)

It can be derived from the definitions that while short-range dependent processes have exponentially decaying autocorrelations, the autocorrelations of long-range dependent processes decay hyperbolically; i.e., the related distribution is heavy-tailed. In practical terms, a random variable with heavy-tail distribution generates extremely large values with high probability [10, 11].

The degree of self-similarity is expressed by the parameter H or Hurst parameter. The parameter represents the speed of decay of a process' autocorrelation function. As $H \to 1$, the extent of both self-similarity and long-range dependence increases. It can also be shown that for self-similar processes with long-range dependency, $H > 0.5$ [12].

1.2. Traffic Models

Traffic modeling originates in traditional voice networks. Most of the models have relied on the assumption that the underlying processes are Markovian [13] (or more general, short-range dependent). However, today's high-speed digital packet networks are more complex and bursty than traditional voice traffic due to the diversity of network services and technologies.

Several sophisticated stochastic models have been developed as a reaction to new developments, such as Markov-modulated Poisson processes [14], fluid flow models [15], Markovian arrival processes [16, 17], batched Markovian arrival process models [19], packet train models [19, 20, 21], and Transform-Expand-Sample models [22]. These models mainly focus on the related queuing problem analytically. They are usually not compared to real traffic patterns and not proved that they match the statistical property of actual traffic data [23, 24, 25, 26].

Another category of models attempts to characterize the statistical properties of actual traffic data. For a long time, the area of networking research has lacked adequate traffic measurements [27]. However, during the past years, large quantities of network traffic measurements have become available and have been collected on the Web and in high-speed networks. Some of these data sets consist of high-resolution traffic measurements over hours, days, or weeks, (e.g., [11, 28, 29, 30, 31, 32, 33, 34, 36]. Other data sets provide information over time periods ranging from weeks to months and years. See, for instance, in [37, 38], signaling system 7 (SS7) data sets in [39, 40] and WAN measurements [11, 28, 40, 41, 42, 43]. Statistical analyses of these high time-resolution traffic measurements have proved that actual traffic data from packet networks reveal self-similarity. These results point out the difference between traditional models and measured traffic data. While the assumed processes in traditional packet traffic models are short-range dependent, measured packet traffic data show evidence of long-range dependency. *Figure 1* illustrates the difference between Internet traffic and voice traffic for different numbers of aggregated users [44]. As the number of voice flows increases, the traffic becomes more and more smoothed contrary to the Internet traffic.

Quite the opposite to the well-developed field of short-range dependent queuing models, there are fewer theoretical results exist for queueing systems with long-range dependence. For some of the results, see [45, 46, 47].

In terms of modeling, the two major groups of self-similar models are fractional Gaussian noises and fractional ARIMA processes. The Gaussian models accurately represent aggregation of many traffic streams as it has been proved in [48] and illustrated by simulation in [49]. The results in [4] demonstrated the effect of long-range dependency on queue length statistics of a single queue system through a controlled fractionally differenced ARIMA (1, d, 0) input process. Another well-known model, the M/Pareto model, has been used in modeling network traffic that is not sufficiently aggregated for the Gaussian model to apply, for example [49, 50].

1.2.1. Black Box versus Structural Models

We share the authors' opinion in [51, 53] calling the approach of traditional time series analysis as black box modeling as opposite to the structural modeling that focus on the environment in which the models' data was collected, i.e., the complex hierarchies of network components that make up today's communications systems. While the authors admit that black box models can be and are useful in other contexts, they argue that black box models are of no use for understanding the dynamic and complex nature of the traffic in modern packet networks. Black box models have not much use in designing, managing and controlling today's networks either. To provide physical explanations for empirically observed phenomena such as long-range dependency, we need to replace black box models with structural models. The attractive feature of structural traffic models is that they take into account the details of the layered architecture of today's networks and can analyze the interrelated network parameters that ultimately determine the performance and operation of a network. Time series models usually handle these details as black boxes [51]. Because actual networks are complex systems, in many cases, black box models assume numerous parameters to represent a real system accurately. For network designers, who are important users of traffic modeling, black box models are not very useful. It is rarely possible to measure or estimate the model's numerous parameters in a complex network environment. For a network designer, a model ought to be simple and meaningful in a particular network. It can relay on actual network measurements, and the result ought to be relevant to the performance and the operation of a real network.

For a long time, traffic models were developed independently of traffic data collected in real networks. These models could not be applied in practical network design. Today, the availability of huge data sets of measured network traffic and the increasing complexity of the underlying network structure emphasize the application of the Ockham' Razer in network modeling. (Ockham's Razor is a principle of the medieval philosopher William Ockham.) According to this principle, modelers should not make more assumptions than the minimum needed. This principle is also called the Principle of Parsimony and motivates all scientific modeling and theory building. It states that modelers should choose the simplest model among a set of otherwise equivalent models of a given phenomenon. In any given model, Ockham's Razor helps modelers to include only those variables that are really needed to explain the phenomenon. Following the principle, model development will become easier, reducing the possibilities for inconsistencies, ambiguities, and redundancies [52].

Structural models are presented, for instance, in [53] and [54]. The papers demonstrate how the self-similar nature of aggregated network traffic of all conversations between hosts explains the details of the traffic dynamics at the level generated by the individual hosts. The papers introduce structural traffic models that have a physical meaning in the network context and underline the predominance of long-range dependence in the packet arrival patterns generated by the individual conversations between hosts. The models provide insight into how individual network connections behave in LANs and WANs. Although the models go beyond the black box modeling methodology by taking into account the physical structure of the aggregated traffic patterns, they do not include the physical structure of the intertwined structure of links, routers, switches, and their finite capacities along the traffic paths.

In [3], the authors demonstrated that World Wide Web (WWW) traffic shows characteristics that are consistent with self-similarity. They show that transmission times may be heavy-tailed due to the distribution of available file sizes in the Web. It is also shown that silent times may also be heavy-tailed, primarily due to the effect of user "think time." Similarly to the structural models in [53, 54], this paper lacks analysis of the impact of self-similar traffic on the parameters of the links and the routers' buffers that ultimately determine a networks' performance.

The effect of self-similar traffic on queue length statistics for real asynchronous transfer mode (ATM) variable bit rate (VBR) traffic is studied in [55]. Making use of fractionally differenced ARIMA (1, d, 0) for the input process, the authors study the queue length process for a varying level of mean utilization. The results of the simulation model demonstrate that the increased buffer size does not significantly decrease cell-loss probabilities for long-range dependent processes. The models developed under the S+ package [56] focus on the queue length of a single-server queue. They do not simulate the combined, overall effect of the interrelated, individual components of a network, such as links, buffers, and applications' response times.

Our paper describes a traffic model that belongs to the aforementioned structural model category. We implement the M/Pareto model within the discrete event simulation package COMNET that allows the analysis of the negative impact of self-similar traffic on not just one single queue, but on the overall performance of various interrelated network components, such as link, buffers, response times, etc. The commercially available package does not readily provide tools for modeling self-similar, long-range dependent network traffic. The model-generated traffic is based on measurements collected from a real ATM network. The choice of the package emphasizes the need for integrated tools that could be useful not just for theoreticians, but also for network engineers and designers. Our paper intends to narrow the gap between existing, well-known theoretical results and their applicability in everyday, practical network analysis and modeling. It is highly desirable that appropriate traffic models should be accessible from measuring, monitoring, and controlling tools. Our model can help network designers and engineers, the ultimate users of traffic modeling, to understand the dynamic nature of network traffic and assist them in designing, measuring, monitoring, and controlling today's complex, high-speed networks in their everyday practice.

1.2.2. Implications of Burstiness on High-Speed Networks
Various papers [33, 50, 57, 58, 59] discuss the impact of burstiness on network congestion. Their conclusions are as follows:

- Congested periods can be quite long with losses that are heavily concentrated.
- Linear increases in buffer size do not result in large decreases in packet drop rates.
- A slight increase in the number of active connections can result in a large increase in the packet loss rate.

Results show that packet traffic "spikes" (which cause actual losses) ride on longer-term "ripples," which in turn ride on still longer-term "swells" [33].

Another area where burstiness can affect network performance is a link with priority scheduling between classes of traffic. In an environment where the higher priority class has no enforced bandwidth limitations (other than the physical bandwidth), interactive traffic might be given priority over bulk-data traffic. If the higher priority class is bursty over long time scales, then the bursts from the higher-priority traffic could obstruct the lower-priority traffic for long periods of time.

The burstiness may also have an impact on networks where the admission control mechanism is based on measurements of recent traffic, rather than on policed traffic parameters of individual connections. Admission control that considers only recent traffic patterns can be misled following a long period of fairly low traffic rates.

2. Model Parameters

Each transaction between a client and a server consists of active periods followed by inactive periods. Transactions consist of groups of packets sent in each direction. Each group of packets is called a burst.

The burstiness of the traffic can be characterized by the following time parameters:

- Transaction Interarrival Time (TIAT): The time between the first packet in a transaction and the first packet of the next immediate transaction.
- Burst Interarrival Time ($1/\lambda$, λ: arrival rate of bursts): The time between bursts.
- Packet Interarrival Time (1/r, r: arrival rate of packets): The time between packets in a burst.

2.1. The Hurst Parameter
It is anticipated that the rapid and ongoing aggregation of more and more traffic onto integrated multiservice networks will eventually result in traffic smoothing. Once the degree of aggregation is sufficient, the process can be modeled by Gaussian process [49]. Currently, network traffic does not show characteristics that close to Gaussian. In many networks, the degree of aggregation is not enough to balance the negative impact of bursty traffic. However, before traffic becomes Gaussian, existing methods can still provide accurate measurement and prediction of bursty traffic.

Most of the methods are based on the estimate of the Hurst parameter H-the higher the value of H, the higher the burstiness and, consequently, the worse the queuing performance of switches and routers along the traffic path. Some are more reliable than others. The reliability depends on several factors, e.g., the estimation technique, sample size, time scale, traffic shaping or policing, etc. Based on published measurements, we investigated methods with the smallest estimation error1. Among those, we chose the Rescaled Adjusted Range (R/S) method because we found it implemented in the Benoit package [60]. The Hurst parameter calculated by the package is input to our method.

2.2. The M/Pareto Traffic Model and the Hurst Parameter
Recent results in [49, 61] have proved that the M/Pareto model is appropriate for modeling long-range dependent traffic flow characterized by long bursts. Originally, the model was introduced in [61] and applied in the analysis of ATM buffer levels. The M/Pareto model was also used in [68, 69, 70, 71] to predict the queuing performance of Ethernet, VBR video, and Internet protocol (IP) packet streams in a single server queue. We apply the M/Pareto model not just for a single queue, but also for predicting the performance of an interconnected system of links, switches, and routers affecting the individual network elements' performance. We make use of some of the calculations presented in [49, 61, 68, 69, 70].

The M/Pareto model is a Poisson process of overlapping bursts with arrival rate λ. A burst generates packets with arrival rate r. Each burst, from the time of its interval, will continue for a Pareto-distributed time period. The use of Pareto distribution results in generating extremely long bursts that characterize long-range dependent traffic.

The probability that a Pareto-distributed random variable X exceeds threshold x is:

$$(1)\quad P(X > x) = \begin{cases} \left(\dfrac{x}{\delta}\right)^{-\gamma}, & x \geq \delta \\ 1, & otherwise \end{cases}$$

$$1 < \gamma < 2, \delta > 0.$$

The mean of X, the mean duration of a burst $\mu = \delta\gamma/(\gamma - 1)$ and its variance is infinite [68]. Assuming a t time interval, the mean number of packets M in the time interval t is:

$$(2)\quad M = \lambda \mathrm{tr}\delta\gamma/(\gamma-1)\text{, and}$$
$$(3)\quad \lambda = M(\gamma-1)/\mathrm{tr}\delta\gamma$$

The M/Pareto model is described in [68, 69, 70, 71] as asymptotically self-similar, and it is shown that for the Hurst parameter, the following equation holds:

$$(4)\quad H = (3-\gamma)/2$$

3. Implementation of the Hurst Parameter in the COMNET Modeling Tool

We implemented the Hurst parameter and a modified version of the M/Pareto model in the discrete event-simulation system COMNET [64, 65]. By using discrete event-simulation methodology, we can get realistic results in measuring network parameters, such as utilization of links and the queuing performance of switches and routers. Our method can model and measure the harmful consequences of aggregated bursty traffic and predict its impact on the overall network's performance.

3.1. Traffic Measurements

To build the baseline model, we collected traffic traces in a large corporate network by the Concord Network Health network analyzer system. We took measurements from various broadband and narrowband links, including 45 Mbps ATM, 56 Kbps, and 128 Kbps frame-relay connections. The Concord Network Health system can measure the traffic in certain time intervals at network nodes, such as routers and switches. We set the time intervals to 6,000 seconds and measured the number of bytes and packets sent and received per second, packet latency, dropped packets, discard eligible packets, etc. Concord Network Health cannot measure the number of packets in a burst and the duration of the bursts, as it is assumed in the M/Pareto model. Due to this limitation of our measuring tool, we slightly modify our traffic model according to the data available. We took snapshots of the traffic in every five minutes from a narrowband frame-relay connection between a remote client workstation and a server at the corporate headquarters as traffic destination in the format in *Table 1*.

The mean number of bytes, the message delay from the client to server, the input buffer level at the client's local router, the number of blocked packets, the mean utilizations of the 56 Kbps frame relay, the digital signal (DS–3) segment of the ATM network, and the 100 Mbps Ethernet link at the destination are summarized in *Table 2*.

COMNET represents a transaction by a message source, a destination, the size of the message, communication devices, and links along the path. The rate at which messages are sent is specified by an interarrival time distribution-the time between two consecutive packets. The Poisson distribution in the M/Pareto model generates bursts or messages with arrival rate λ—the number of arrivals, which are likely to occur in a certain time interval. In simulation, this information is expressed by the time interval between successive arrivals $1/\lambda$. For this purpose, we use the Exponential distribution. Using the Exponential distribution for interarrival time will result in an arrival pattern characterized by the Poisson distribution. In COMNET, we implemented the interarrival time with the function Exp $(1/\lambda)$. The interarrival time in the model is set to one second matching the sampling time interval set in Concord Network Health and corresponding to an arrival rate λ = 1/sec.

In the M/Pareto model, each burst continues for a Pareto-distributed time period. The Concord Network Health can-

TABLE 1

Traffic Traces

Delta Time (Sec)	Average Bandwidth Utilization %	Bytes Total/Sec	Bytes In/Sec	Bytes Out/Sec
299	2.1	297	159.2	137.8
300	2.2	310.3	157.3	153
300	2.1	296.8	164.4	132.4
300	2.7	373.2	204.7	168.5
……..	……..	……..	….	….

TABLE 2

Measured Network Parameters

Mean # of Bytes Total/Sec	Message Delay from Client to Server in ms	Input Buffer Level at Client's Router in Bytes	Dropped Packets	*Links' Mean Bandwidth Utilization%*		
				56 Kbps Frame Relay	DS–3 in ATM Network	100 Mbps Ethernet at Destination
440.4279	78.687	0.04	0	3.14603	0.06	0.0031

FIGURE 2

Part of the Real Network Topology Where the Measurements Were Taken

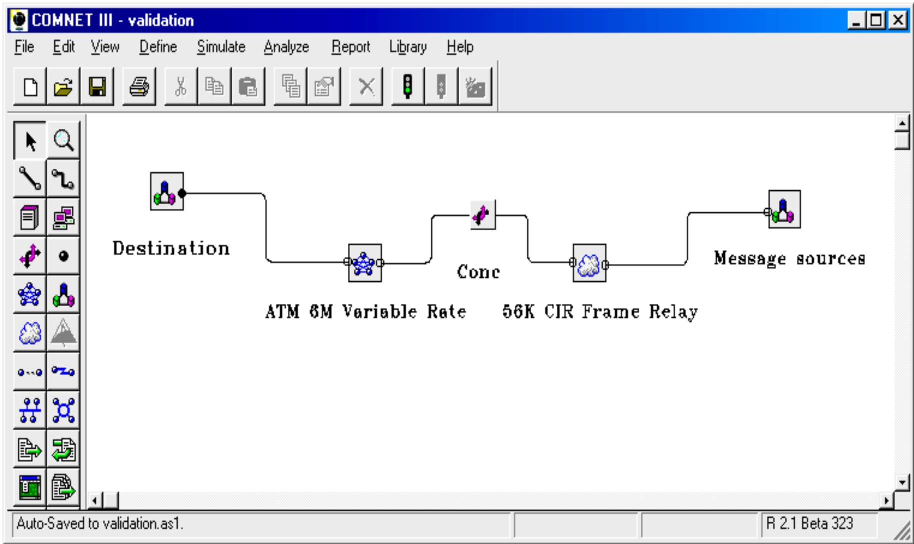

not measure the duration of a burst; hence, we assume that a burst is characterized by the number of bytes in a message sent or received in a second. Since the ATM cell-rate algorithm ensures that equal length messages are processed in equal time, longer messages then require longer processing time. So we can say that the distribution of the duration of bursts is the same as the distribution of the length of bursts. Hence, we can modify the M/Pareto model by substituting the Pareto-distributed duration of bursts with the Pareto-distributed length of bursts. We derive δ of the Pareto distribution not from the mean duration of bursts, but rather from the mean length of bursts.

The Pareto distributed length of bursts is defined in COMNET by two parameters: the location and the shape. The location parameter corresponds to the δ, the shape parameter corresponds to the γ parameter of the M/Pareto model in (1) and can be calculated from the relation (4) as

(5) $\gamma = 3 - 2H$

The Pareto distribution can have infinite mean and variance. If the shape parameter is greater than 2, both the mean and variance are finite. If the shape parameter is greater than 1, but less than or equal to 2, the mean is finite, but then the variance is infinite. If the shape parameter is less than or equal to 1, both the mean and variance are infinite.

From the mean of the Pareto distribution we get the following:

(6) $\delta = \mu * (\gamma - 1)/\gamma$

The relations (5) and (6) allow us to model bursty traffic based on real traffic traces by performing the following steps:

a. Collect traffic traces using the Concord Network Health network analyzer.
b. Compute the Hurst parameter H by making use of the Benoit package with the traffic trace as input.
c. Use the Exponential and Pareto distributions in the COMNET modeling tool with the parameters calculated as previously described to specify the distribution of the interarrival time and length of messages.
d. Generate traffic according to the modified M/Pareto model and measure network performance parameters.

The traffic generated according to these steps is bursty, with parameter H calculated from real network traffic.

4. Validation of the Baseline Model

We validate our baseline model by comparing various model parameters of a 56 Kbps frame-relay and a 6 Mbps ATM connection with the same parameters of a real network as the Concord Network Health network analyzer traced it. For simplicity, we use only the "Bytes Total/Sec" column of the trace, i.e., the total number of bytes in the "Bytes Total/Sec" column is sent in one direction only from the client to the server. The Hurst parameter of the real traffic trace is H = 0.55 calculated by the Benoit package. The topology is shown in *Figure* 2.

The "Message sources" icon is a subnetwork that represents a site with a token-ring network, a local router, and a client A sending messages to the server B in the "Destination" subnetwork (see *Figure* 3).

The interarrival time and the length of messages (see *Figure* 4) are defined by the Exponential and Pareto functions Exp (1) and Par (208.42, 1.9) respectively. The Pareto distribution's location (208.42) and shape (1.9) are calculated from formulas (5) and (6) by substituting the mean length of bursts (440 bytes from *Table* 2) and H = 0.55.

The corresponding heavy-tailed Pareto probability distribution and cumulative distribution functions are illustrated in the *Figure 5*. (The X-axis represents the number of bytes.)

The "Frame Relay" icon represents a frame-relay cloud with 56K committed information rate (CIR). The "Conc" router connects the frame-relay network to a 6 Mbps ATM network with variable rate control (VBR), as shown in *Figures 6* and *7*.

FIGURE 3

"Message Source" Remote Client

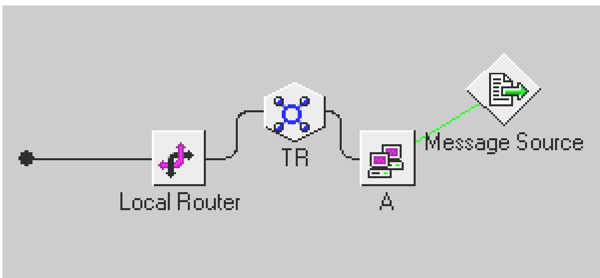

FIGURE 4

Interarrival Time and Length of Messages Sent by the Remote Client

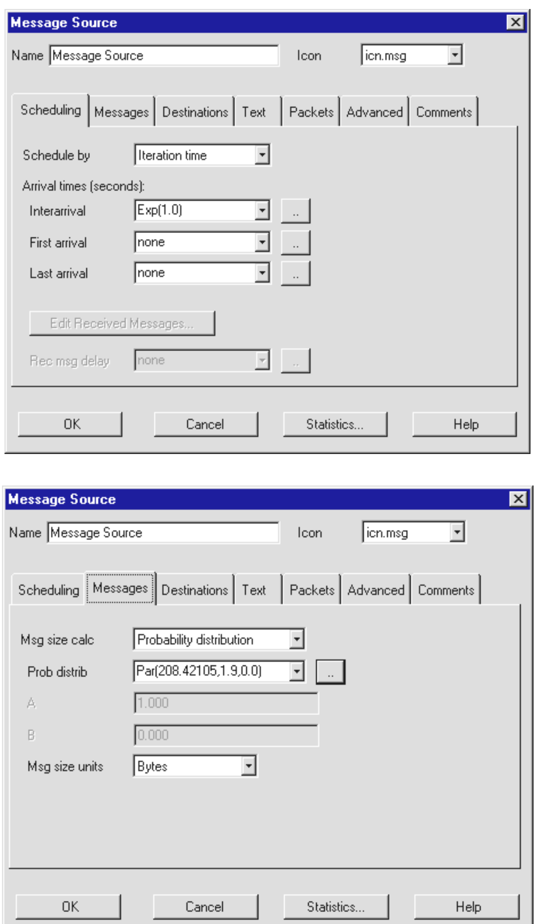

The "Destination" icon denotes a subnetwork with server B (see *Figure 8*).

The results of the model show almost identical average utilization of the frame-relay link (0.035 ~ 3.5%) to the utilization of the real measurements (3.1%) (see *Figure 9*).

The message delay in the model is also very close to the measured delay between the client and the server (78 msec) (see *Figure 10*).

The input buffer level of the remote client's router in the model is almost identical with the measured buffer level of the corresponding router (see *Figure 11*).

Similarly, the utilizations of the model's DS–3 link segment of the ATM network and the Ethernet link in the destination network closely match with the measurements of the real network (see *Figure 12*).

It can also be shown from the model's traffic trace that for the model-generated messages, the Hurst parameter H = 0.56, i.e., the model generates almost the same bursty traffic as the real network.

Furthermore, the number of dropped packets in the model was zero, similar to the number of dropped packets in the real measurements. Therefore, we start from a model that closely represents the real network.

5. Consequences of Traffic Burstiness

To illustrate our method, we developed a COMNET simulation model to measure the consequences of bursty traffic on network links, message delays, routers' input buffers, and the number of dropped packets due to the aggregated traffic of large number of users. The model implements the Hurst parameter as it has been described in Section 3. We repeated the simulation for 6,000 seconds, 1,6000 seconds, and 18,000 seconds to allow infrequent events to occur a reasonable number of times. We found that the results are very similar in each simulation.

FIGURE 5

The Pareto Probability Distribution for Mean 440 Bytes and Hurst Parameter H = 0.55

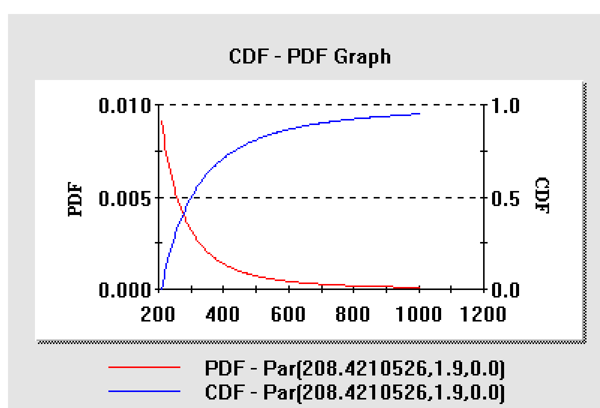

FIGURE 6

The Internal Links of the 6 Mbps ATM Network with Variable Rate Control (VBR)

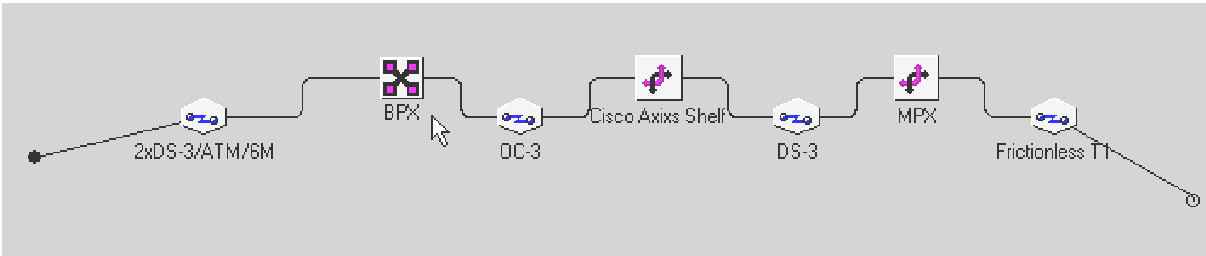

5.1. Topology of Bursty Traffic Sources

The "Message Source" subnetworks transmit messages as in the aforementioned baseline model, but with different burstiness: H = 0.95, H = 0.75, H = 0.55, and with fixed size (see *Figure 13*). Initially, we simulate four subnetworks and four users per subnetwork, each sending the same volume of data (mean 440 bytes per second) as in the validating model.

FIGURE 7

Parameters of the 6Mbps ATM Connection

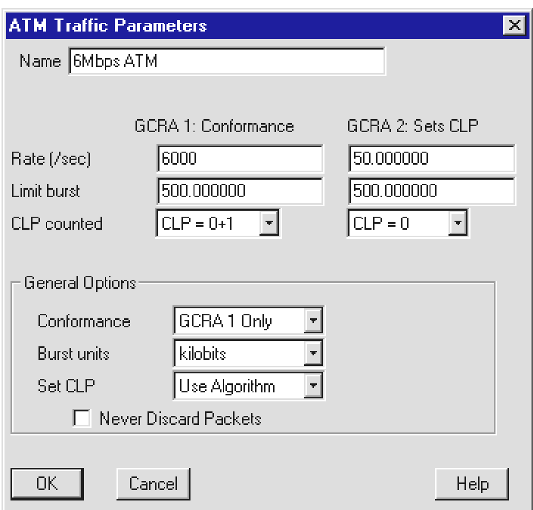

5.2. Link Utilization and Message Delay

First, we are going to measure and illustrate the extremely high peaks in frame-relay link utilization and message delay. The model traffic is generated with message sizes determined by various Hurst parameters and fixed-size messages for comparison. The COMNET modeling tool has a trace option to capture its own model-generated traffic. It has been verified that for the model-generated traffic flows with various Hurst parameters, the Benoit package computed similar Hurst parameters for the captured traces.

Table 3 shows the simulated average and peak link utilization of the different cases. The utilization is expressed in the [0,1] scale, not in percentages:

The charts in the Appendix clearly demonstrate that even though the average link utilization is almost identical, the frequency and the size of the peaks increase with the burstiness, causing cell drops in routers and switches. We received the following results for response time measurements (see *Table 4*).

The charts in the Appendix graphically illustrate the relation between response times and various Hurst parameters.

5.3. Other Network Performance Parameters

In addition to the utilization of the individual frame-relay links carrying traffic with increasing burstiness, we can also

FIGURE 8

The "Destination" Subnetwork

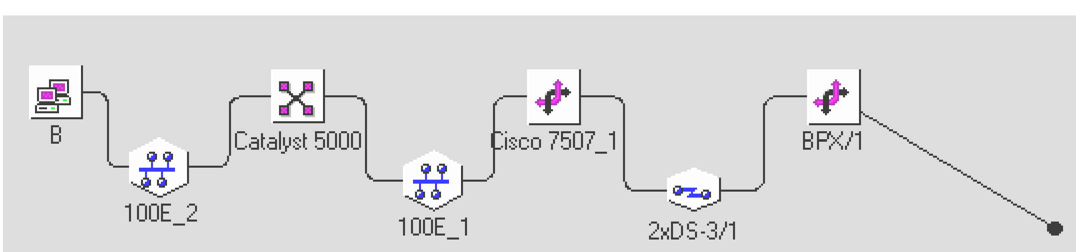

FIGURE 9

Utilization of the Frame-Relay Link in the Baseline Model

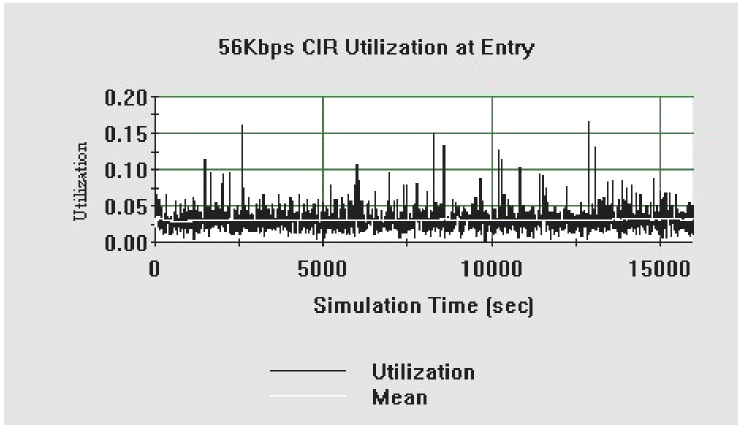

FIGURE 10

Baseline Message Delay between the Remote Client and the Server

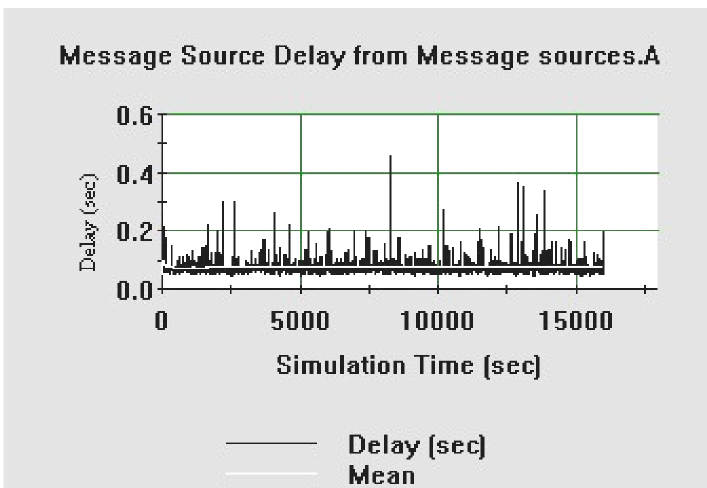

FIGURE 11

Input Buffer Level of Remote Router

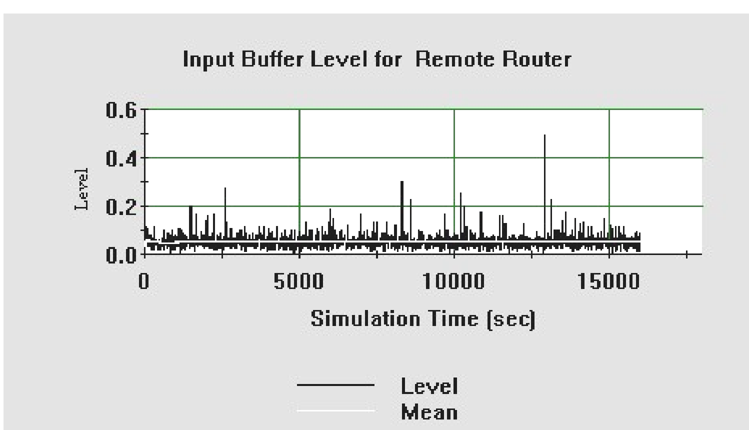

FIGURE 12

Baseline Utilizations of the DS-3 Link and Ethernet Link in the Destination

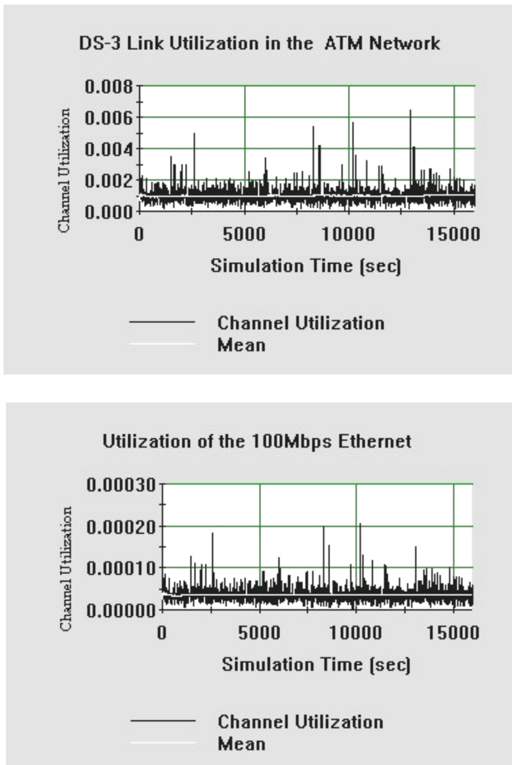

measure the impact of the combined traffic on various network parameters of the ATM network and the final destination network, such as link utilizations and input buffer levels. The charts in the Appendix depict these simulation results as well.

5.4. Input Buffer Level for Large Number of Users

We also measured the number of cells dropped at a router's input buffer in the ATM network due to a surge of bursty cells. We simulated the aggregated traffic of approximately 600 users, each sending the same number of bytes in a second as in the measured real network. The number of blocked packets is summarized in *Table 5*.

6. Conclusion

The paper presented a discrete event-simulation methodology to measure various network performance parameters while transmitting bursty traffic. It has been proved in recent studies that combining bursty data streams will also produce bursty combined data flow. The studies imply that the methods and models used in traditional network design require modifications. We categorize our modeling methodology as a structural model [51, 53] contrary to a black box model. Structural models focus on the environment in which the models' data was collected, i.e., the complex hierarchies of network components that make up today's communications systems. Although black box models are useful in other contexts, they are not easy to use in designing, managing, and controlling today's networks. We implemented a well-known model, the M/Pareto model, within the discrete event-simulation package COMNET that allows the analysis

FIGURE 13

Network Topology of Bursty Traffic Sources with Various Hurst Parameters

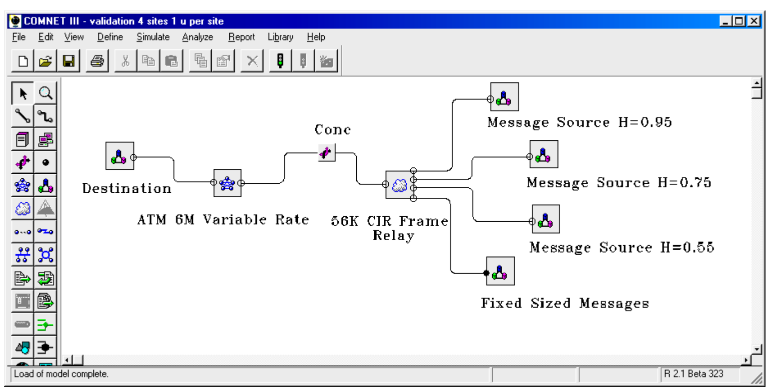

TABLE 3

Simulated Average and Peak Link Utilization

	Fixed-Size Messages	H = 0.55	H = 0.75	H = 0.95
Average Utilization	0.12	0.13	0.13	0.14
Peak Utilization	0.18	0.48	1	1

TABLE 4

Response Time and Burstiness

	Fixed-Size Messages	H = 0.55	H = 0.75	H = 0.95
Average Response Time (ms)	75.960	65.61	87.880	311.553
Peak Response Time (ms)	110.06	3510.9	32418.7	112458.08
Standard Deviation	0.470	75.471	716.080	4341.24

TABLE 5

Relation between the Number of Cells Dropped and Burstiness

	Fixed-Size Messages	H = 0.55	H = 0.75	H = 0.95
Packets Accepted	13282	12038	12068	12622
Packets Blocked	1687	3146	3369	7250
Average Buffer Use in Bytes	56000858	61001835	62058222	763510495

of the negative impact of self-similar traffic on not just one single queue, but on the overall performance of various interrelated network components as well. Using real network traces, we built and validated a model by which we could measure and graphically illustrate the impact of bursty traffic on link utilization, message delays, and buffer performance of frame-relay and ATM networks. We illustrated that increasing burstiness result in extremely high link utilization, response time, and dropped packets and measured the various performance parameters by simulation.

The choice of the package emphasizes the need for integrated tools that could be useful not just for theoreticians, but also for network engineers and designers. Our paper intends to narrow the gap between existing, well-known theoretical results and their applicability in everyday, practical network analysis and modeling. It is highly desirable that appropriate traffic models should be accessible from measuring, monitoring, and controlling tools. Our model can help network designers and engineers, the ultimate users of traffic modeling, to understand the dynamic nature of network traffic and assist them in their everyday practice.

Appendix

Measurements for Link Utilization

The following charts (see *Figures 14* through *17*) demonstrate that even though the average link utilization for the various Hurst parameters is almost identical, the frequency and the size of the peaks increase with the burstiness causing cell drops in routers and switches. The utilization is expressed in the [0,1] scale not in percentages.

FIGURE 14

Utilization of the Frame-Relay Link for Fixed-Size Messages

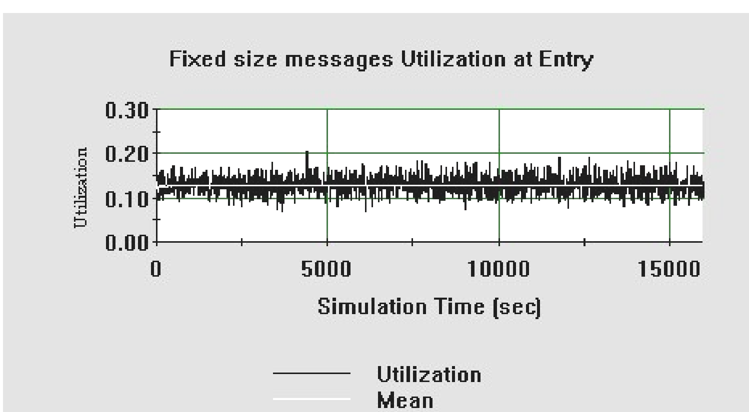

FIGURE 15

Utilization of the Frame-Relay Link for Hurst Parameter H = 0.55

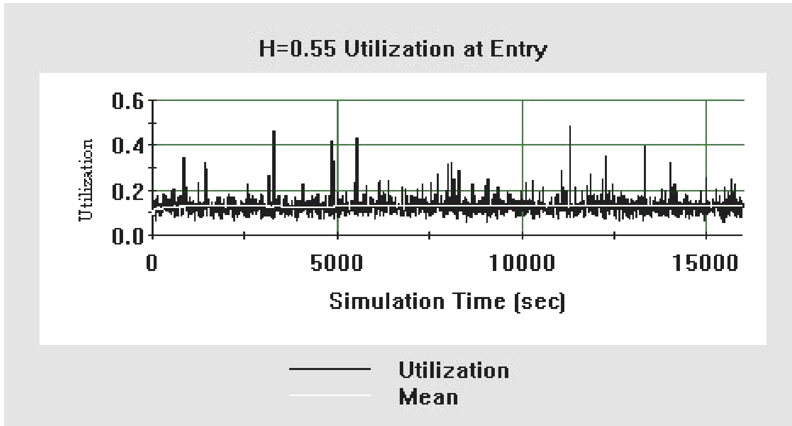

FIGURE 16

Utilization of the Frame-Relay Link for Hurst Parameter H = 0.75 (Peaks Are Higher)

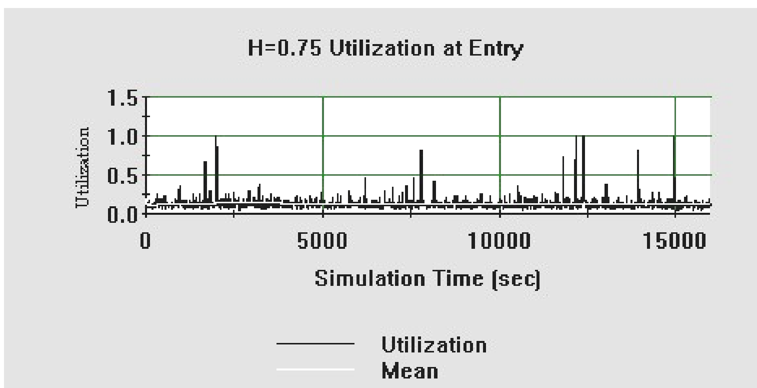

FIGURE 17

Utilization of the Frame-Relay Link for Hurst Parameter H = 0.95 (Many High Peaks)

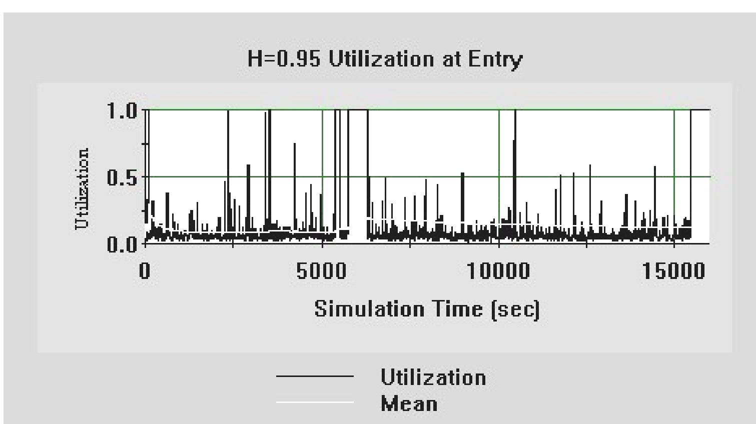

Measurements for Message Delays

The following charts (see *Figures 18* through *21*) illustrate the relation between response time and various Hurst parameters.

FIGURE 18

Message Delay for Fixed-Size Messages

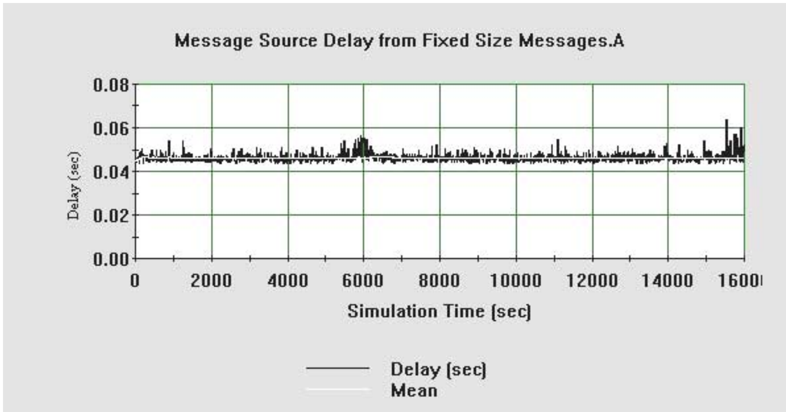

FIGURE 19

Message Delay for H = 0.55 (Longer Response-Time Peaks)

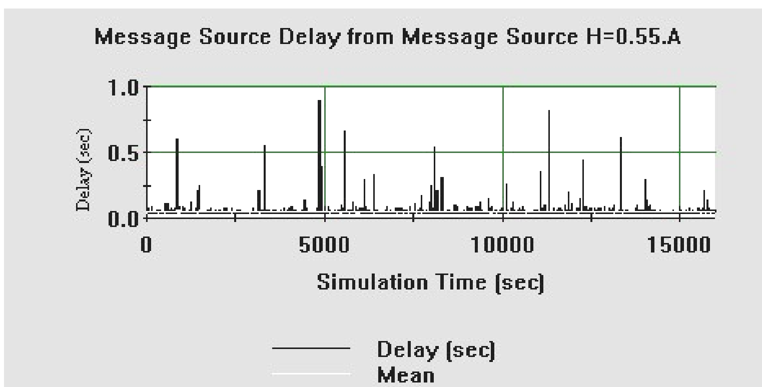

FIGURE 20

Message Delay for H = 0.75 (Long Response-Time Peaks)

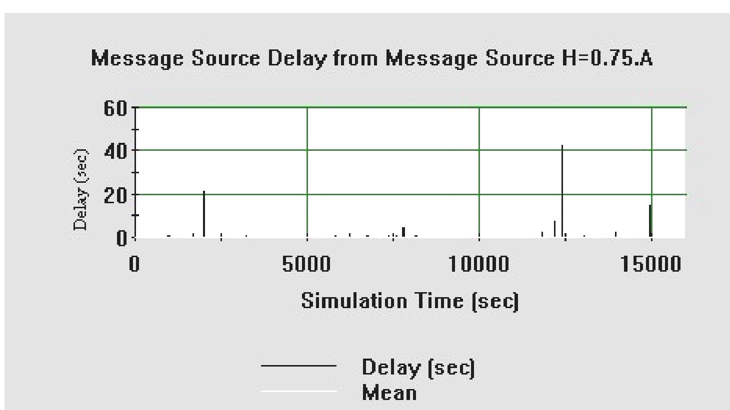

FIGURE 21

Message Delay for H = 0.95 (Extremely Long Response-Time Peak)

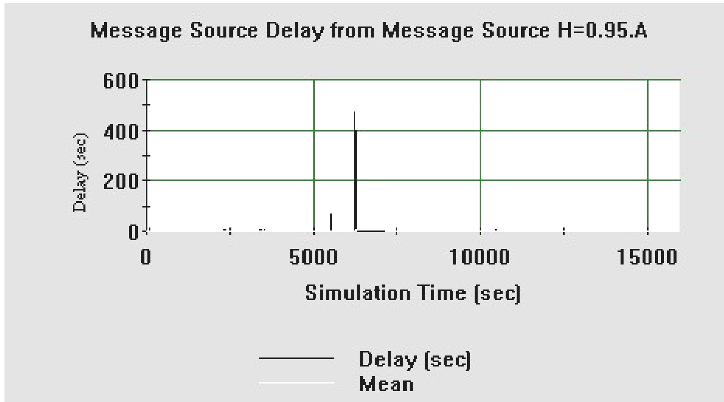

Measurements of Other Network Performance Parameters

We also measured the impact of the combined traffic on various network parameters of the ATM network and the final destination network, such as the DS–3 link utilizations, the 100 Mbps Ethernet, and input buffer levels of the Cisco-7500 and the BPX switches. The following charts (see *Figures 22* through *25*) depict these simulation results.

FIGURE 22

Utilization of the DS–3 link in the 6 Mbps ATM Network

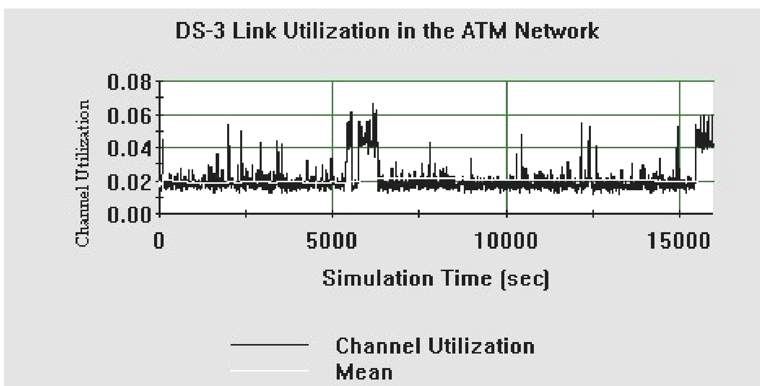

FIGURE 23

Utilization of the 100 Mbps Ethernet Link in the Destination Network

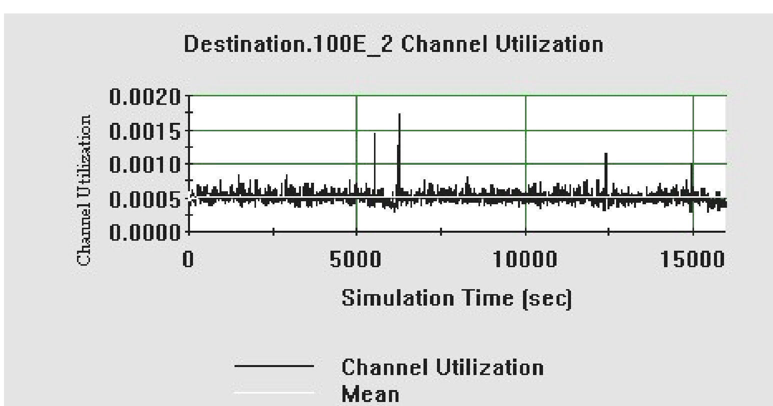

FIGURE 24

Input Buffer Level of the Cisco 7507 Router

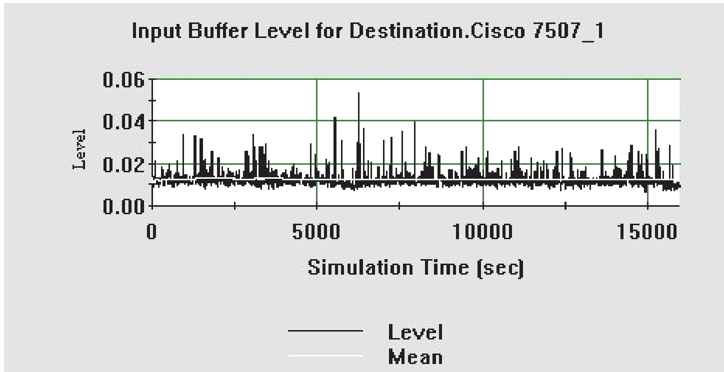

FIGURE 25

Input Buffer Level of the BPX Switch

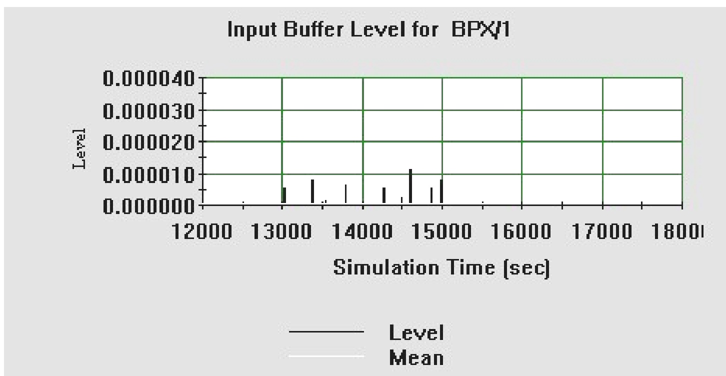

References

[1] W. Leland et al., "On the self-similar nature of Ethernet traffic," Extended Version, IEEE/ACM Transactions on Networking, Vol. 2, No. 1 (1994), 1-15.

[2] W. Leland et al., "Statistical analysis and stochastic modeling of self-similar data traffic," in The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks, eds. J. Labetoulle and J. W. Roberts, Proceedings of the 14th International Teletraffic Congress (ITC '94) (Amsterdam: Elsevier Science B.V., 1994), 319-328.

[3] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," IEEE/ACM Transactions on Networking, Vol. 5., No. 6. (December 1997).

[4] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queuing analysis with long-range dependent packet traffic," IEEE/ACM Transactions on Networking, Vol. 4, No. 2 (April, 1996), 209-223.

[5] W. Leland and D. Wilson, "High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection," Proceedings of IEEE NFOCOM'91 (Bal Harbour, Fla.: 1991), 1,360-1,366.

[6] K. Park, G. Kim, and M. Crovella, "On the relationship between file sizes, transport protocols, and self-similar network traffic," Proceedings of the 4th Int. Conf. Network Protocols (ICNP'96) (October 1996), 171-180.

[7] W. Willinger et al., "Self-similarity in high-speed packet traffic: Analysis and modeling of Ethernet traffic measurements," Statist. Sci., Vol. 10, No. 1 (1995), 67-85.

[8] J. Beran, "Statistics for Long-Memory Processes," Monographs on Statistics and Applied Probability (London: Chapman and Hall, 1994).

[9] D. Cox, "Long-Range Dependence: A Review," in Statistics: An Appraisal, H. A. David and H. T. David (Eds.) (Ames, Iowa: The Iowa State University Press, 1984), 55-74.

[10] B. Mandelbrot, The Fractal Geometry of Nature (San Francisco, Calif.: W.H. Freeman and Co., 1982).

[11] V. Paxson and S. Floyd, "Wide-area traffic: The failure of Poisson modeling," IEEE/ACM Transactions on Networking, Vol. 3 (June 1995), 226-244.

[12] B. Mandelbrot, J. Van Ness, "Fractional Brownian Motions, Fractional Noises and Applications", SIAM Review 10 (1968), 422-437.

[13] W. Willinger and M. Taqqu, "Bibliographical Guide to Self-Similar Traffic and Performance Modeling for Modern High-Speed Networks," in Stochastic Networks: Theory and Applications, eds. F. Kelly, S. Zachary and I. Ziedins (Oxford: Clarendon Press, 1996), 339-366.

[14] H. Hefles and D. Lucantoni, "A Markov modulated characterization of packetized voice and data traffc and related statistical multiplexer performance," IEEE Journal on Selected Areas in Communications, 4 (1986), 856-868.

[15] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a datahandling system with multiple sources," Bell System Technical Journal, 61 (1982), 1,871-1,894.

[16] M. F. Neuts, "A versatile Markovian point process," Journal of Applied Probability, 18 (1979), 764-1,779.

[17] M. F. Neuts, "Structured Stochastic Matrices of M/G/1 Type and Their Applications," (New York: Marcel Dekker, 1989).

[18] S. Kogon and D. Manolakis, "Linear parametric models for signals with long-range dependence and infinite variance," ICASSP'95 (1995).

[19] R. Jain and S. Routhier, "Packet trains: Measurements and a new model for computer network traffic," IEEE Journal on Selected Areas in Communications, 4 (1986), 986-995.

[20] V. Ramaswami and G. Latouche, "Modeling packet arrivals from asynchronous input lines," Proceedings of the 12th ITC (Torino, Italy: 1988).

[21] V. Ramaswami et al., "Comparison of some traffic models for ATM performance studies," in Teletraffic and Datatraffic in a Period of Change,

eds. A. Jensen and V. B. Iversen, Proceedings of ITC-13 in Copenhagen, Denmark (North Holland, Amsterdam: 1991), 7-12.

[22] B. Melamed, "An overview of TES processes and modeling methodology," in Models and techniques for Performance Evaluation of Computer and Communications Systems, eds. L. Donatiello and R. Nelson (New York: Springer-Verlag, 1993), 359-393.

[23] A. Arvidsson and R. Harris, "On the performance of some commonly employed models of bursty traffic," Preprint (1994).

[24] C. Partridge, "The end of simple traffic models," Editor's Note, IEEE Network (September 1993).

[25] V. Ramaswami, "Traffic performance modeling for packet communications: Whence where and whither?" Proceedings of the 3rd Australian Teletraffic Seminar (Adelaide, Australia: 1988).

[26] W. Willinger et al., "On traffic measurements that defy traffic models (and vice versa): self-similar traffic modeling for high-speed networks," ConneXions, 8(11) (1994), 14-24.

[27] P. Pawlita, "Two decades of data traffic measurements: A survey of published results, experiences and applicability," in Teletraffic Science for New Cost-Effective Systems, Networks and Services, ed. M. Bonatti, ITC-12 (North Holland: Elsevier Science Publishers B.V., 1989).

[28] J. Beran et al., "Long-range dependence in variable-bit-rate video traffic," IEEE Transactions on Communications, 43 (1995), 1,566-1,579.

[29] R. Caceres et al., "Characteristics of wide-area TCP/IP conversations," Proceedings of the ACM Sigcomm'91 (1991), 101-112.

[30] D. Duffy et al., "Analyzing telecommunications traffic data from working common channel signaling subnetworks," in Statistical Applications of Expanding Computer Facilities, Volume 25, eds. M. Tarter and M. Lock (Interface Foundation of North America, 1993), 156-165.

[31] M. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," Proceedings of the ACM/SIGCOMM'94 (London, UK: 1994), 269-280.

[32] R. Gusella, "A measurement study of diskless workstation traffic on an Ethernet," IEEE Transactions on Communications, 38 (1990), 1,557-1,568.

[33] W. Leland et al., "On the self-similar nature of Ethernet traffic," Computer Communications Review, 23, Proceedings of the ACM/SIGCOMM'93 (San Francisco, Calif.: September 1993), 183-193.

[34] W. Leland and D. Wilson, "High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection," Proceedings of the Infocom '91 (Bal Harbour, Fla.: 1991), 1,360-1,366.

[35] T. Neame, M. Zukerman, and R. Addie, "Applying Multiplexing Characterization to VBR Video Traffic," Proceedings of ITC 16 (June 1999).

[36] K. E. E. Raatikainen, "Symptoms of self-similarity in measured arrival process of Ethernet packets to a file server," Technical Report Series of Publications C-1994-4; Dept. of Computer Science; University of Helsinki; 1994.

[37] V. Paxson, "Growth trends in wide-area TCP connections," IEEE Network (July/August 1994), 8-17.

[38] K. Claffy, H. Braun, and G. Polyzos, "Tracking long-term growth of the NSFNET backbone," Proceedings Infocom '93 (San Francisco, Calif.: 1993).

[39] D. Duffy et al., "Analyzing telecommunications traffic data from working common channel signaling subnetworks," in Statistical Applications of Expanding Computer Facilities, Volume 25, eds. M. E. Tarter and M. D. Lock (Interface Foundation of North America, 1993), 156-165.

[40] D. Duffy et al., "Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks," IEEE Journal on Selected Areas in Communications, 12 (1994), 544-551.

[41] K. Claffy, H. Braun, and G. Polyzos, "Application of sampling methodologies to network traffic characterization," Proceedings of the ACM/SIGCOMM '93 (San Francisco, Calif.: 1993), 194-203.

[42] V. Paxson, "Growth trends in wide-area TCP connections," IEEE Network (July/August 1994), 8-17.

[43] R. Gibbens, " The statistical analysis of broadband traffic," in Stochastic Networks: Theory and Applications, eds. F. Kelly, S. Zachary, and I. Ziedins (Oxford: Clarendon Press, 1996), 339-366.

[44] M. Listanti and V. Eramo, "Architectural and Technological Issues for Future Optical Internet Networks," IEEE Communications (September 2000), 82-92.

[45] N. Duffield, "Economies of scale in queues with sources having power-law large deviation scalings," DIAS-APG-94-27, Dublin Institute of Advanced Studies; Dublin, Ireland; 1994.

[46] N. Duffield and N. O'Connell, "Large deviation and overflow probabilities for the general single-server queue, with applications," Mathematical Proceedings of the Cambridge Philosophical Society, 118 (1995), 363-375.

[47] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queuing analysis with long-range dependent packet traffic'" IEEE/ACM Transactions on Networking, 4 (1996), 209-223.

[48] R. G. Addie, "On the weak convergence of long range dependent traffic processes," submitted to J. Stat. Planning and Inference.

[49] R. Addie, M. Zukerman, and T. Neame, "Broadband Traffic Modeling: Simple Solutions to Hard Problems," IEEE Communications (August 1998), 88-95.

[50] Neame et al., "Investigation of traffic models for high speed data networks," Proceedings of ATNAC '95 (1995).

[51] W. Willinger and V. Paxson, "Discussion of "Heavy Tail Modeling and Teletraffic Data by S.R. Resnick," The Annals of Statistics, 25(5) (1997), 1,805-1,869.

[52] Francis Heylighen, http://pespmc1.vub.ac.be/HEYL.html.

[53] W. Willinger, V. Paxson, and M.Taqqu, "Self-similarity and heavy tails: Structural modeling of network traffic," in A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tailed Distributions, eds. R. Adler, R. Feldman, and M. S. Taqqu (Boston: Birkhauser Verlag, 1998).

[54] W. Willinger et al., "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level," Extended Version, IEEE/ACM Transactions on Networking, 5 (1997), 71-86.

[55] A. Adas and A. Mukherjee, "On Resource Management and QoS guarantees for Long Range Dependent Traffic," Preprint (1994).

[56] The S+ package, Version 3.0 (Statistical Sciences, Inc., September 1991).

[57] P. J. Davy, "Self-similar processes with non-negative stationary increments," submitted to J. of Stat. Planning and Inference.

[58] S. Molnar, A. Vidacs, and A. Nilsson, "Bottlenecks on the Way towards Fractal Characterization of Network Traffic: Estimation and Interpretation of the Hurst Parameter," http://hsnlab.ttt.bme.hu/~molnar.

[59] M. Taqqu, V. Teverovsky, and W. Willinger, "Estimators for long-range dependence: an empirical study," Fractals, Vol. 3, No. 4 (1995), 785-788.

[60] Benoit 1.1, Trusoft Int'l Inc., www.trusoft-international.com.

[61] N. Likhanov, B. Tsybakov, and N. Georganas, "Analysis of an ATM Buffer with Self-Similar ('Fractal') Input Traffic," Proceedings of the IEEE INFOCOM Conference (Boston: 1995), 985-992.

[62] Gyires, T., "Methodology for Modeling the Impact of Traffic Burstiness on High-Speed Networks," Proceedings of the 1999 IEEE International Conference on Systems, MAN, and Cybernetics (Tokyo: October 12-15, 1999), 980-985.

[63] Communications with COMNET developers Vipul Pabari and Ken Neumeister (1999).

[64] "COMNET III, Planning for Network Managers," Release 1.4, 1997, CACI Products Company, 333 North Torrey Pines Ct., La Jolla, California, 92037.

[65] "Enterprise Profiler," CACI Products Company, 1998, 3333 North Torrey Pines Ct., La Jolla, California, 92037.

[66] Gyires, T., "Smart Permanent Virtual Circuits in Asynchronous Transfer Mode Networks," Invited Paper in the Proceedings of the 1997 IEEE International Conference on Systems, MAN, and Cybernetics (Orlando, Fla.: October 12-15, 1997), 4,460-4,464.

[67] Gyires, T., "Performance Prediction of Smart Permanent Virtual Circuits in ATM Networks with CACI COMNET," Proceedings of the 1998 IEEE International Conference on Systems, MAN, and Cybernetics (San Diego, Calif.: October 11-14, 1998), 3,993-3,998.

[68] T. Neame, M. Zukerman, and R. Addie, "A Practical Approach for Multi-Media Traffic Modeling," Proceedings of Broadband Communications '99 (Hong Kong: November 10-12, 1999), 73-82.

[69] T. Neame and M. Zukerman, "Application of the M/Pareto Process to Modeling Broadband Traffic Streams," Proceedings of ICON '99 (Brisbane, Queensland, Australia: September 28 - October 1, 1999), 53-58.

[70] T. Neame, M. Zukerman, and R. Addie, "Modeling Broadband Traffic Streams," Proceedings of Globecom '99 (Rio de Janeiro, Brazil: December5-9, 1999).

[71] T. Neame, M. Zukerman, and R. Addie, "Applying Multiplexing Characterization to VBR Video Trafficc," Proceedings of ITC 16 (June 1999).

Notes

1. Variance, Aggregated Variance, Higuchi, Variance of Residuals, Rescaled Adjusted Range (R/S), Whittle Estimator, Periodogram, Residuals of Regression [59].

Efficient Regression Testing of CTI Systems: Testing a Complex Call-Center Solution

Andreas Hagerer
Senior Engineer
METAFrame Technologies GmbH

Tiziana Margaria
Managing Partner and Chief Executive Officer
METAFrame Technologies GmbH

Oliver Niese
Chair of Programming Systems, Department of Computer Science
University of Dortmund

Bernhard Steffen
Chair of Programming Systems, Department of Computer Science
University of Dortmund

Georg Brune
Director, Development and Support for Applications and Solutions
Siemens AG

Hans-Dieter Ide
Head of HO SE 1
Siemens AG

Abstract

In this paper, we show how the workfllow-oriented, library-based test-design and execution environment presented in [8] fulfills the promise of enabling efficient integrated system-level tests of complex industrial applications. We exemplify on the basis of a concrete case study (Siemens' HPCO Application, a complex call-center solution) how test engineers can now work with the integrated test environment (ITE). Efficiency measurements in the field document an average performance gain in the execution of functional regression tests of factors beyond 30: Setting up the test scenario is now the most expensive activity!

1. Introduction

The world of telecommunications has rapidly evolved during the last 15 years, modifying in this process its focus. In 1985, a telephone switch was used "only" as a telephone switch. Additional components, either hardware or software, were gradually developed to bring additional functionality and flexibility to the traditional switch, e.g., in the initial days, voice-mail or billing systems. Today, not only are single functionalities added at a quick pace, but the switch is also mutating its role into the central element of complex heterogeneous and multivendor systems: It is nowadays integrated into whole business solutions, e.g., in the field of hotel solutions, call centers, and unified messaging applications. Technically, whereas in the earlier days the interaction between the switch and the applications was almost exclusively implemented via proprietary interfaces, today the definition of open standards—such as the computer-supported telecommunication applications (CSTA) protocol [1,2], which specifies the command sets and the data structures needed at the application layer, or the telephony application programming interface (TAPI) [6], which defines the programming interface for telephony applications—pushes the field toward the development of new, system-level applications.

Figure 1 (left Y axis) documents the trend toward increased product integration in terms of the augmentation in the number of value-added applications that work on average with or on a switch. As one can see, the integration factor has been rapidly increasing since 1995, and this trend points in the direction of even larger value-added product ranges.

A parallel but concurring aspect is the increasing number of major switch releases per year (see *Figure 1*, right Y axis). This trend is driven mainly by the convergence between the classical telecommunications technology and the modern Internet protocol (IP) technology, i.e., voice over IP. Modern

FIGURE 1

Trends in the Development of CTI Systems—(left Y axis) Growing Product Integration and (right Y axis) Faster-Paced Switch Releases

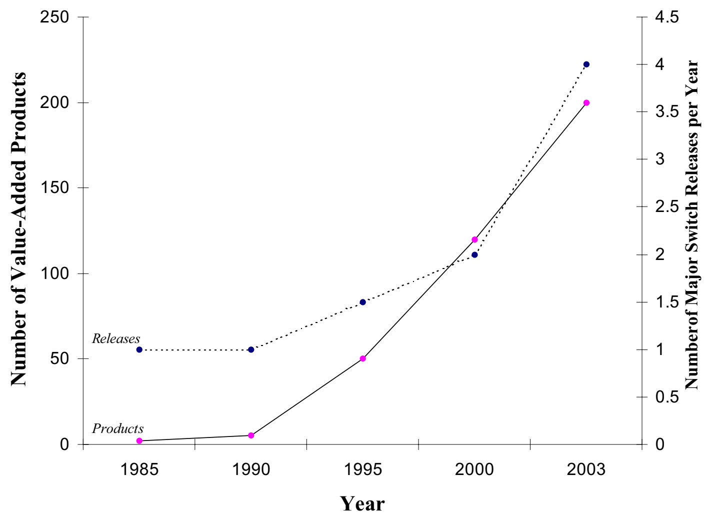

switches are themselves complete complex systems, and they experience the accelerating evolution pace of hardware and software in combination!

In this rapidly evolving scenario, the need for efficient automated regression testing is evident. Whenever a release arises—either of the switch or of (a subset of) the application programs that cooperate with it, and either singularly or, increasingly often, in collaborative combinations—the correct functioning of the new configurations must be certified again.

Altogether, testing complex telephony solutions is a multidimensional task that demands automation via adequate system-level tool support. The complexity lies in the interaction between the components as well as in the short innovation cycles and the great number of possible combinations between the private branch exchange (PBX) and the value-added applications. Therefore, an adequate environment must focus on structuring, efficiency, and abstraction.

In summary, systems-under-test have become composite (e.g., including computer telephony integrated [CTI] platform aspects), embedded (due to hardware/software codesign practices), reactive, and run on distributed architectures (e.g., client/server architectures). Complex subsystems affect each other in a variety of complex ways, so mastering today's testing scenarios for telephony systems demands an integrated, open, and flexible approach to support the management of the overall test process, i.e., specification of tests, execution of tests, and analysis of test results.

In [8], we presented an automated testing environment for CTI systems capable of handling their structural complexity. Our approach offers a coarse-grained testing environment, realized in terms of a component-based test design on top of a library of elementary but intuitively understandable test-case fragments. The relations between the fragments are treated orthogonally, delivering a test design and execution environment enhanced by means of lightweight formal verification methods.

Now, one year later, we summarize our practical experiences in the use of this environment at Siemens, where the ITE has been used in a production environment. In this paper, we show on the basis of a concrete case study (Siemens' HPCO Application) how test engineers now work with the ITE. Measurements in the field have exceeded our expectations, documenting an average efficiency improvement in the test execution of functional regressions of complex call-center solutions and of several other CTI applications of factors beyond 30: As mentioned, setting up the test scenario is now the most expensive activity.

In the rest of the paper, we first describe in Section 2 the concrete industrial application (regression test of a complex call-center solution), then we detail in Section 3 the process followed to integrate the call-center solution into the ITE. Section 4 presents our field results, concerning in particular the concrete measured speed-up of the test execution, and Section 5 imparts our conclusions and some perspectives.

2. Regression Test of a Complex Call-Center Solution

A typical example of an integrated CTI platform is a call-center solution, as illustrated in *Figure 2*, where the switch is connected to the integrated services digital network (ISDN) telephone network, or, more generally, to the public switched telephone network (PSTN), and acts as a "normal" telephone switch to the phones. Additionally, it communicates directly via a local-area network (LAN) or indirectly via an application server with call-center client/server applications that are executed on personal computers (PCs). Like the phones, these CTI applications are active components; they may stimulate the switch (i.e., initiate calls), and they also react to stimuli sent by the switch (i.e., notify incoming calls). Therefore, in a system-

FIGURE 2

Test Setting for a Complex Call-Center Solution

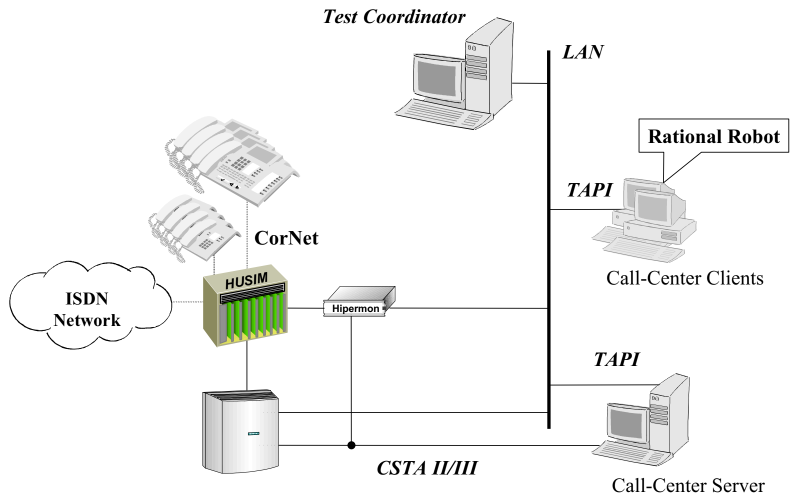

level test scenario, it is necessary to investigate the interaction between such subsystems.

Even the relatively simple scenario of *Figure 2* demonstrates the complexity of CTI platforms from the communications point of view, because there are several (internal) protocols involved. For example, the telephones communicate via the corporate network protocol[1] with the PBX, whereas the PBX communicates via the CSTA Phase II/III protocol [1,2] with the application server. On the application server, a TAPI service provider performs a mapping of the CSTA protocol to the TAPI protocol [6], which is the communications protocol between the application server and its clients.

2.1. The Integrated Test Environment

The ITE (see [7, 8] for a more detailed discussion) is based on an existing general-purpose environment for the management of complex workflows, METAFrame Technologies' Agent Building Center (ABC) [11], which contains built-in features concerning test coordination and test organization. The test coordinator constitutes the test-management layer of our environment and includes an application-specific specialization of the ABC for the domain of system-level regression testing of telephony systems.

To communicate with different test tools, a flexible common object request broker architecture (CORBA)/remote method invocation (RMI)–based architecture has been developed. The test coordinator executes integrated test cases by controlling several test tools, each managing its own subsystem. The extensibility of the environment by additional test tools is the key of the approach.

3. Integration of the Call-Center Solution into the ITE: The Process

To be able to handle regression testing of the call-center solution within the ITE, the following steps have to be performed:

- Integration of test tools
- Identification and implementation of testblocks
- Test-case design and execution

These steps, which must be taken whenever a new system-under-test is integrated into the ITE, are discussed in detail in the following subsections.

3.1. Integrating the Test Tools

3.1.1. How to Integrate

The ITE offers a generic CORBA interface to communicate with the different test tools that comprise the basic methods (e.g., getName, getVersion) that all test tools have to support. Specific features (e.g., input and output commands for testing a special component) can be added as need by extending the interface. The specific methods of the test tools are accessed via test-tool–specific basic functionalities (testblocks).

The integration process consists of two main activities:

1. Integration of the interface provided by the test tool into the test coordinator
2. Implementation of the interface functionality by the test tool

If the interface is not implemented in the test tool by the vendor, there are different ways for integration.

Plug-In Approach: If the test tool supports customization via loading plug-ins or libraries, the interface can be integrated by implementing such a plug-in, including the CORBA object request broker (ORB).

Separate Server Process: If a test tool offers remote access via an interface of its own (e.g., COM, DCOM, or CORBA), a separate server process can communicate via this interface with the test tool. Then, the special server implements the interface or its derivation and communicates with the test coordinator—i.e., it can be seen as a relay between the test coordinator and the test tool.

3.1.2. What to Integrate
Two different test tools are used in the call-center scenario:

1. A proprietary protocol analyzer (Hipermon [5]), which is connected to a telephone simulator (Husim) and to the connection between the switch and the application server
2. A graphical user interface (GUI) test tool (Rational Robot [10]), which is used in several instances—i.e., for every considered call-center client

Hipermon/Husim: The required interface to the ITE is implemented in the Hipermon/Husim directly by the vendor. It provides the tester with full access to the whole functionality of a telephone as well as to the underlying protocols, i.e., corporate network (CorNet) protocol and CSTA II/III. The Hipermon is able to decode CorNet/CSTA messages and to deliver the attached data to the ITE for further processing.

Rational Robot: The Rational Robot does not provide a CORBA implementation itself. Here we followed the plug-in approach: We implemented an external library that provides the CORBA functionality and brings the Rational Robot into an interpreter mode, which allows us to send Robot Scripts to the Rational Robot for execution. Furthermore, it is possible to use generalized scripts, which use variables where the concrete values are requested at run time from the ITE.

3.2. Identification of Suitable Testblocks
Several classes of testblocks had to be provided to test the call-center solution:

- *Testblocks for the User Interfaces:* The user-interface layer concerns the functionality of the features that constitute the user interfaces of the single subsystems, such as telephone devices (e.g., the hook switch, display, and lamps of a telephone set) and software applications (e.g., the buttons, check boxes, and text fields of an application's GUI).
- *Testblocks for Communication between the Switch and Telephone Devices*: Devices, both simulated and real, communicate with the switch according to the CorNet protocol.
- *Testblocks for Communication between the Switch and Call-Center Server*: The CSTA protocol [1, 2] specifies the command sets and the data structures needed for this application layer.

3.2.1. Defining Testblocks for the User-Interface Layer
Generally, tests of GUIs are performed by means of "recording and replaying" tools, which capture events and effects and allow storage of them for later execution. The Rational Robot [10] has been made available in the ITE in a flexible way for driving an application's GUI. A testblock is generated automatically for each Robot Script, which is, in turn, either recorded or programmed. As shown in *Figure 3* (top left), the code of a Robot Script is enriched by special keywords that

FIGURE 3

Automatic Generation of Test Building Blocks

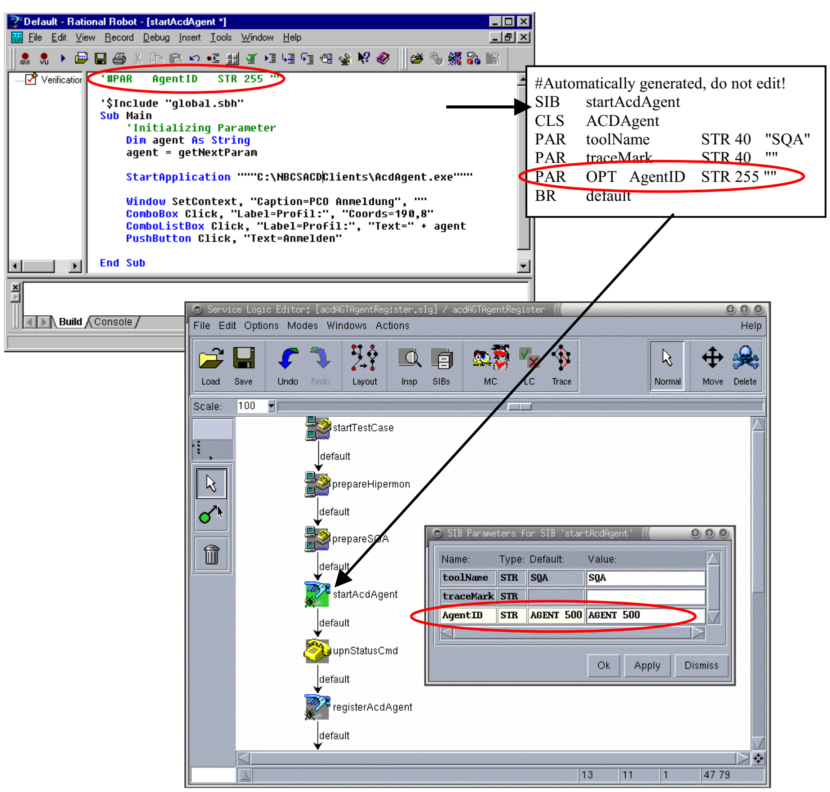

specify the parameters of the script. These parameters are essential for the test coordinator in order to instruct the Rational Robot during the test execution.

This solution enabled the test engineers at Siemens to provide more than 80 different test blocks for a call center (i.e., log-on/log-off of an agent) in a short time by simply recording the actions on the GUI.

3.2.2. Defining Testblocks for the CorNet Communication Layer
Messages exchanged between switch and device according to the CorNet protocol:

- Are used to modify the display lines, brightness, and color of lamps or the ringer mode and ringer pattern.
- Initiate calls, indicate incoming calls, or terminate calls.
- Activate or deactivate specific features of the switch.

Such messages are collected by the Hipermon, which decodes them into a format appropriate for a further evaluation within the ITE.

In this context, testblocks are needed that either test the current state of a device or simulate the functionality of a device toward the switch.

3.2.3. Defining Testblocks for the CSTA Communication Layer
CSTA defines a telephony process model for applications and a computing process model for the switch. An implementation of CSTA consists of CSTA services and the CSTA-protocol. By means of the protocol, an application accesses the CSTA telephony services from the switch or provides CSTA computing services to the switch. Each model consists of a set of objects and rules to change the states of the objects. Examples of telephony objects include the following:

- *A device object*, representing anything that allows users to access telecommunications services. It can be either a physical device (buttons, lines, and stations) or logical devices (a group of devices, an automatic call distribution [ACD] group). A device has attributes, including *device type, device identifier*, and *device state* that can be monitored and manipulated by an application.

  In the CTI system of *Figure 3*, a special CSTA device (agent) is used. An agent represents the association and the activities of a physical device with an ACD group. An agent—i.e., the physical device—becomes associated with a specific ACD group by a process of logging on. The agent's state describes the current relation of the agent with the ACD group (logged-on, ready to accept calls, busy, working after call).

- *A call object*, describing the logical session among calling and called parties. The call behavior—i.e., its establishment and release—can be observed and manipulated by an application. A call object representing a call session has attributes such as *identifier* and *state* and offers operations such as *make* or *clear*. One or more devices may be involved in a call in different phases.

- *A connection object*, representing a relationship between a call and a device. It is characterized by attributes such as *identifier* and *state* and by operations such as *hold* or *clear*. Many CSTA services, such as hold-call, reconnect-call, or clear-call, are made by operating on connections comprising a call.

Basically, CSTA services consist of a request and a response, both possibly parameterized. A monitor service can be activated to track control and other activities and to receive notification of all changes of the switch. Starting a monitor indicates that an application, be it a component of the system or an external observer, wants to be notified of changes that occur in calls, devices, or applications as well as device attributes managed by the switch. Examples of changes include arrivals of a call at a device, answering a call, and changing a device by modifying features such as "forwarding." Event reports are sent to the monitor-requestor. For the test of CTI systems it is necessary to have access to the CSTA events for validation.

3.2.4. How to Identify and Implement Testblocks
To design system-level test cases, it is necessary to know which features and functions the system provides, how to operate the system in order to stimulate a feature, and how to determine whether the features work. This information is gathered for each key component of the system and each functional area, resulting in a set of stimulation actions and verification actions that form the basis for the design of test cases in terms of so-called testblocks. These are specified by means of the following three steps:

1. *Classification:* A testblock name is identified and organized relative to a hierarchical classification scheme. In the call-center application, we have classes for each subsystem and each layer.
2. *Parameterization:* A set of formal parameters is defined to enable a more general usage of a testblock.
3. *Connectivity:* To steer the control flow during test execution, each testblock possesses a set of outgoing branches representing the different possibilities for continuation.

3.3. Design and Execution of Testgraphs

Once the libraries of testblocks are available, we are ready to design executable test cases. As described in detail in [8] and shown in a very simple case in [9], designing test cases consists of a behavior-oriented graphical combination of testblocks. Icons representing testblocks are graphically stuck together to yield testgraph structures that embody the test behavior in terms of control. A macro technique allows the encapsulation of a section of a testgraph into a new testblock to support a hierarchical design of test cases. Parameters of testblocks that are used in a macro can be specified as parameters of the macro.

The design of test cases is controlled by validation techniques that enforce essential technical frame conditions and guarantee the key features of the underlying test purpose, as illustrated in detail in [8]. There we showed how critical consistency requirements, including version compatibility and frame conditions for executability, are formulated and how consistency of test cases is fully automatically enforced via model checking and error diagnosis.

In the ITE, testgraphs are immediately executable by means of ABC's tracer module. Starting at a specified testblock of a testgraph, the tracer proceeds from testblock to testblock. The actions associated with each testblock are performed, i.e., stimuli and inspection requests are sent to the corresponding system's component, and the responses are received and evaluated. The evaluation results successively determine the selection of the next testblock.

4. Evaluation

To evaluate the economic impact of the ITE introduction, we must first identify the cost factors that pertain to testing CTI systems. *Table 1* identifies the macroscopic cost factors that arise along the test-process life cycle together with their frequency of occurrence.

Test planning and specification, and the definition and set-up of test scripts, occur only initially, when an experimental scenario (i.e., the testing of a specific CTI system) is set up. The planning and specification phases are not (yet) affected by the ITE. The usual collection or programming of test scripts, which in a manual setting directly constitute the elementary testblocks, is in the ITE supported by a largely automated generation of reusable testblocks that fit with the overall ITE architecture.

The main focus of the ITE in this first phase is the reduction of costs for the repetitive phases of CTI testing. Primarily, we addressed test execution, the former bottleneck (see *Table 2*), in combination with the automatic creation of test reports and (in the near future) advanced support of the analysis of test results.

Table 2 documents the measured improvement of the test-execution costs due to the introduction of the ITE. The systems under test considered in each row are composed by the PC client/server application listed in column 1, which cooperates with the HICOM switch along the configuration pattern illustrated in *Figure 2*. The second and third column report the measured effort (in man hours) of one regression cycle for the system under test when performed manually (column 2) or with the ITE (column 3). The improvement is dramatic: speed-up factors of about 30 for each regression-cycle execution! Indeed, the personnel is now completely freed from manual test execution; the test cases run in batch mode, e.g., at night or while the test engineers work on further test-case definitions or on setting up a different test scenario. The automated protocol facility allows them to just look off-line at the test cases that ended up with errors or warnings, providing, in this case, detailed information.

The shift toward an automated test-execution environment brings additional benefits that go beyond the pure savings in manpower. In particular, the test-design activities during the test planning and specification phases have shifted from largely programming to a more conceptual modeling and generation approach. This has effects on the organization and the maintainability of the test suites (there is a high-level model of the test cases), on the documentation of the test runs (if all goes well, few data are stored), on their repeatability (the configurations, etc., are well documented and automatically managed), and on the error diagnosis, where the testgraph provides an intuitive guide when looking for causes of misbehavior.

The full automation of test execution is for the moment not yet feasible, since some manual steps—such as system set-up and configuration (e.g., physical connection of the components, installation of the software on the machines)—are still needed. While in a manual execution, setting these steps represented about 20 percent of the effort for a regression cycle.

TABLE 1

Regression Test Cost Factors

Task	Without ITE	With ITE	Frequency
Test Planning	Manual	Manual	Once
Test Specification	Manual	Manual	Once
Test Scripts	Manual	Manual	Once
Test Execution	Manual	Automated (Batch)	Recurrent
Test Protocol	Manual	Automated	Recurrent
Test Analysis	Manual	(Selected Cases)	Recurrent

TABLE 2

Test-Execution Effort in Hours per Regression

System-under-Test	Manual	With ITE	Speed-Up Factor
Hotel Solutions	10.0	0.5	20
Call-Center Solutions	43.0	1.0	43
Analog Voice-Mail	23.0	0.5	46
Digital Voice-Mail	20.0	0.5	40
Call-Charge Computer	19.0	0.5	38
Total	115.0	3.0	ca. 38

Now, due to the test-execution speed-up introduced by the ITE, they are responsible for more than 80 percent of the total effort! We are therefore currently aiming at reducing (or, ideally, automating) the configuration process [4].

5. Conclusion

In this paper, we have shown on the basis of a concrete case study (Siemens' HPCO Application) how test engineers can now work with ITE in their daily regression testing activities. The workfllow-oriented, library-based test-design and execution environment presented in [8] has proved to fulfill the promise of enabling efficient integrated system-level tests of complex industrial applications in real-life usage. The field experience gathered over the last year has shown that the use of ITE has brought a new dimension of efficiency (with measured speed-up factors of greater than 30 for the test-execution phase) for regression testing of large CTI systems together with an improved test-design methodology and test-process organization. A global cost-benefit calculation thus shows that the additional investment for the ITE is well able to pay off in a short period of time if extensively adopted. The ITE, in fact, dramatically reduces the recurring cost factors almost without impairing the remaining positions that concern the basic effort that still has to be spent along the whole test life cycle (test planning, manual configuration of the test settings, etc.) and the necessary upfront investments (e.g., license fees for test tools, hardware, etc.). In fact, we believe that we will be able to also cut down on the required basic effort in the near future.

6. References

1. European Computer Manufactures Association, "Services for Computer Supported Telecommunications Applications (CSTA) Phase II," ECMA 217/218 (1994).
2. European Computer Manufactures Association, "Services for Computer Supported Telecommunications Applications (CSTA) Phase III," ECMA 265/2285 (1998).
3. S. Gladstone, *Testing Computer Telephony Systems and Networks* (Telecom Books, 1996).
4. A. Hagerer et al., "Model Generation by Moderated Regular Extrapolation," to appear in proceedings of the FASE 2002 Int. Conference on Fundamental Approaches to Software Engineering (Grenoble, France: 2002).
5. Herakom GmbH, http://www.herakom.de.
6. Microsoft Coorperation, "Using TAPI 2.0 and Windows to Create the Next Generation of Computer-Telephony Integration," white paper.
7. O. Niese et al., "Library-Based Design and Consistency Checking of System-Level Industrial Test Cases," Proceedings of the FASE 2001 Int. Conference on Fundamental Approaches to Software Engineering, LNCS 2029 (Genova, Italy: Springer Verlag, 2001), 233–248.
8. O. Niese et al., "An Automated Testing Environment for CTI Systems Using Concepts for Specification and Verification of Workflows," *Annual Review of Communications* (Chicago, Ill.: International Engineering Consortium, 2000).
9. O. Niese et al., "Automated Regression Testing of CTI Systems," Proceedings of the IEEE European Test Workshop (ETW 2001) (Stockholm, Sweden: 2001).
10. Rational, Inc., "The Rational Suite Description," http://www.rational.com/products.
11. B. Steffen and T. Margaria, "METAFrame in Practice: Intelligent Network Service Design," *Correct System Design – Issues, Methods and Perspectives*, LNCS 1710, (Springer Verlag, 1999), 390–415.

Notes

1. ECMA and CCITT Q.930/931–oriented D-channel Layer-3 protocol for private IPABX.

Recovery Techniques in Generalized Multiprotocol Label Switching

Robert Haim
Senior Product-Marketing Manager
Polaris Networks, Inc.

Abstract

Generalized multiprotocol label switching (GMPLS) is the evolution of the multiprotocol label switching (MPLS) suite of protocols that extends the control-plane operation used in packet-based networks to optical networks. The basic functions of GMPLS include expediting traffic forwarding on traffic-engineered end-to-end paths (label switched paths, or LSP), creating differentiated services (DiffServ), efficient utilization of the discreet optical bandwidth levels (e.g., optical carrier [OC]–48/192), and dynamic fault-management of LSPs or tunnels. In essence, GMPLS is considered as the common control plane for both electrical and optical networks.

A major factor in the acceptance and successful deployment of GMPLS is its ability to carry out intelligent fault-management. Not only does GMPLS need to provide a more granular level for protection, it must address performance levels that are comparable to SONET protection/restoration schemes (50 ms or better).

While GMPLS uses the synchronous optical network (SONET) protection scheme as a benchmark, it does not have to be bound to the strict "all or nothing" scheme of SONET that gives away 50 percent of available bandwidth for protection. In this paper, GMPLS' varying degrees of protection and restoration schemes are explained.

Brief Primer on GMPLS

The suite of protocols that comprises MPLS is used to speed up packet forwarding and to provide for traffic engineering in Internet protocol (IP) networks. Therefore, the connectionless operation of IP networks become more like a connection-oriented network, where the path between the source and the destination is pre-calculated based on the application's required quality of service (QoS) level. To speed up the forwarding scheme, an MPLS device uses *labels* rather than address matching to determine the next hop for a received packet. To provide traffic engineering, tables are used that represent the levels of QoS that the network can support. The tables and the labels are used together to establish an end-to-end path called a *label switched path* (LSP). Extensions to traditional IP routing protocols (open shortest path first [OSPF], intermediate system–to–intermediate system [IS–IS]) and to the existing signaling protocol, i.e., resource reservation protocol (RSVP), and/or the new signaling protocol, i.e., constraint-based routing–label distribution protocol (CR–LDP), comprise the suite of MPLS protocols.

GMPLS extends MPLS suite of protocols to provide the control plane (signaling, routing, and management) for devices that switch in packet, time, wavelength, and fiber domains. This common control plane promises to simplify network operation and management by automating end-to-end provisioning of connections, managing network resources, and providing the level of QoS that is expected in the new sophisticated applications. To provide a variety of QoS levels that is required by different applications, GMPLS protocols also address network fault-management. For example, link management protocol (LMP) is a new protocol that includes provisions for fault-management in GMPLS.

High-Availability Models for GMPLS

The basic premise of the GMPLS suite of protocols can be viewed as the convergence of control-plane operations for both electrical and optical switching devices. This implies that an end-to-end path (LSP) can be established through different types of networks to carry a variety of services with differentiated levels of QoS. To date, packet switching devices have carried data on a best-effort basis, while SONET–based optical networks have carried highly reliable voice traffic. However, recent applications such as voice traffic over IP, virtual private networks (VPNs), and other multimedia applications have modified the reliability requirements necessary for end-to-end connections. As these types of traffic travel through different optical segments, the protection levels that are available today are not granular enough to satisfy these applications efficiently.

Figure 1 shows a three-segment GMPLS network. The access network(s) support traffic from end-users, such as enterprises, or multi-tenant dwellings, and pass their aggregated multiservice traffic on to metro SONET transport rings. The end users have different QoS–level requirements depending on their service-level agreements (SLAs) with their "access" providers. The equipment in the metro, such as add/drop multiplexers (ADMs) and digital cross-connect systems (DCSs), groom and switch the traffic to the long-haul network, which will further transport the traffic, e.g., via wavelength-based transmission equipment. The deployment of GMPLS in each network segment provides the establishment of end-to-end traffic-engineered LSPs that span all of these segments. And more importantly, it is capable to provide multiple levels of reliability in protecting and restoring services that are specific per end-user application.

SONET protection schemes have been available for some time now and have been proven to provide highly reliable time division multiplexing (TDM) networks that carry voice traffic. These schemes stand as benchmarks for GMPLS. SONET protection schemes are reviews in *Table 1*. We will examine how GMPLS addresses these protection/restoration schemes.

High Availability in GMPLS

Survivability refers to the network's ability to absorb and process deterministic and non-deterministic *changes* while maintaining the level of service that was promised for an application. These *changes* affect network topology, either through planned expansions or unplanned, unforeseen problems. Topological changes and, therefore, network-capacity degradation could also occur when a card, a node, or a line goes down for a variety of reasons. To accommodate network changes in a timely manner, the equipment must have built-in intelligence to assure uninterrupted traffic flow, e.g., via automatic network discovery and dynamic new routes to provide multiple levels of protection and restoration. This intelligence allows seamless operation of the network and bypasses its problems to continue its service. GMPLS–based protection and restoration schemes attempt to provide this intelligence while taking into account that LSPs may span multiple dissimilar network types, e.g., packet or TDM/SONET. Therefore, the main goal for GMPLS–driven reliability is, for a given LSP, which carries a certain QoS–level based on the application's SLA, to identify an *efficient* approach that will maintain the promised QoS based on two "failure" probabilities:

P_{nf} = Probability of nodal failure
P_{lf} = Probability of link failure

For an LSP that spans N nodes and N-1 links, the following can be deduced:

$$P_{lspf} = 1 - (1 - P_{nf})^{N} * (1 - P_{lf})^{N-1}$$

where P_{lspf} is the probability of LSP failure. For simplicity, the formula assumes uniform failure probability for every node and every link that comprise the network. For example, assuming that $P_{nf} = 1 \times 10^{-5}$ and $P_{lf} = 1 \times 10^{-9}$ for a 10 node network, P_{lspf} becomes about 1×10^{-4}. Naturally, P_{lspf} increases and lowers the availability of an LSP, with each additional and/or link. To maintain a given level of QoS in a GMPLS–based network, the following issues have been considered in the industry:

- Must consider that an LSP may traverse a variety of equipment in which the times of LSP set-up may vary greatly from each other, e.g., a high-speed router and a digital cross-connect. The former does not take as much time as the latter does to find an output port for incoming data on an input port.
- GMPLS operates in a multilayer environment where packet-based data is transported by optical cross-connects (OXCs) that are in turn transported by wavelength division multiplexers (WDMs) on a fiber layer. This requires that the layer that is closest to where a

FIGURE 1

End-to-End GMPLS–Based Networks

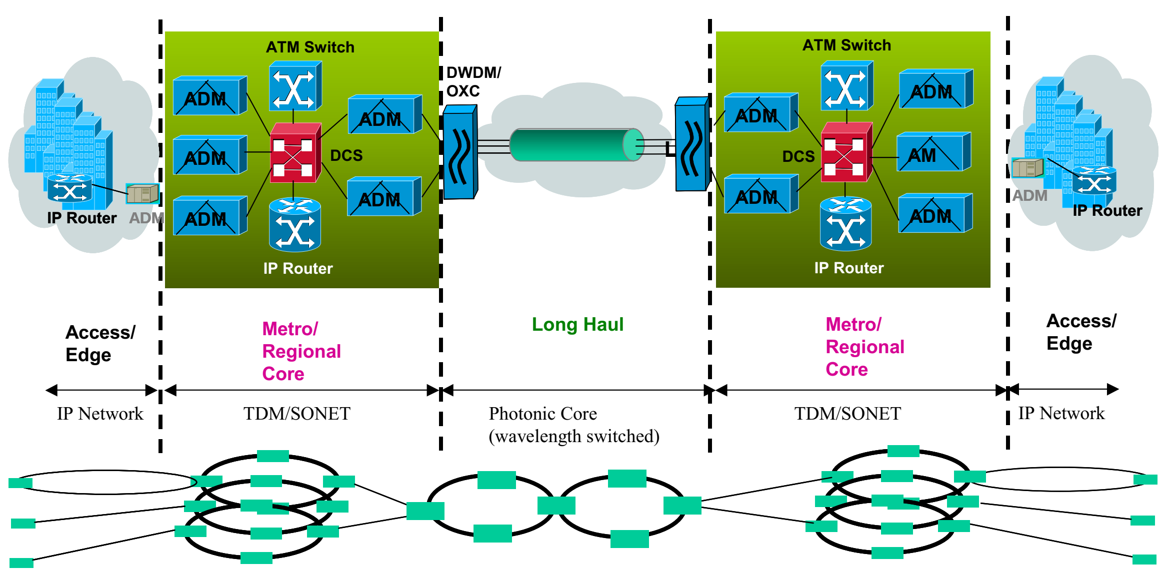

TABLE 1

SONET Protection Schemes

Protection Scheme	Architecture	Level of Protection	Mode of Operation	Usually Used in	Restoration Time
Automatic Protection Switching (APS)	Point to point; only two-line terminating equipment; [Protection provided at line level]	1+1	Signals travel on both working and protection fibers	Access/feeder rings	50 ms or better
		1:N; 1:1	Allows other interruptible traffic to flow on the backup		
Unidirectional Path-Switched Ring (UPSR)	Ring: two fibers [Protection provided at the SONET STS–N path level]	1+1	Does not allow for interruptible traffic to flow on the backup lines	Access rings	50 ms
Bidirectional Line-Switched Ring (BLSR)	Ring/mesh: four fibers [Protection provided at the SONET STS–N path level]	Effectively, 1:1	During normal operation, traffic flows through the four-fiber BLSR on the *working fiber pair*. The *protection fiber pair* is either idle or carrying interruptible traffic during normal operation.	Metro, regional, or long-haul networks	50 ms

fault has occurred must be able to detect it. A recovery action in one layer should not trigger recovery action in another layer.

- Extra, spare capacity must be built into the network but must be used efficiently, i.e., be able to allocate resources, e.g., bandwidth, for protection in such a way that it can be used to carry other traffic when the required QoS allows it.
- Must be able to assign two or more different routes for the same LSP, having end-to-end protection when required by its QoS level.
- For those services that require the highest availability, SONET's 50 ms target time is a sound goal. A good example is a mission-critical application such as real-time distance medical imaging or financial transactions.
- If voice traffic is to be carried on a GMPLS–based LSP, users' call-drop tolerance time of less than two seconds must be taken into account.
- Must also be able to support applications with lower levels of QoS requirements, e.g., e-mail or Web browsing, which require best effort to other real-time services, such as enterprise VPNs, which may require higher levels of QoS than best effort.

GMPLS not only must be capable of providing protection mechanism that are similar to SONET, but it also must be flexible enough to provide more granular protection methods for those applications that do not require it. In the next section of this paper, different protection and restoration schemes that are currently proposed in GMPLS are discussed. For completeness, a brief explanation of protection and restoration is given, since it is important to distinguish between the two schemes.

Protection/Restoration Schemes

The main differences between these two schemes can be summarized as *protection* having a faster operation but being less efficient in use of resources, while *restoration* is slower but uses resources more efficiently. Protection switching requires pre-allocated resources similar to the SONET schemes given in *Table 1*. This means that alternate path(s) must be pre-calculated and established (with or without signaling). Restoration, on the other hand, refers to dynamic allocation of resources, i.e., signaling and establishment of backup paths after the occurrence and detection of fault. Both of these schemes are used in GMPLS to offer different QoS levels to satisfy different application requirements.

Basically, two major linear protection schemes are used: 1+1 and M:N. This is the same as discussed in the SONET protection schemes in *Table 1*. In the former, data is simultaneously transmitted via two distinct paths, called *working* and *protection*. The receiver then autonomously decides which of the two carries the healthier signal. In the 1+1 scheme, the backup path always transports a mirror image of the primary path data while the primary path is operational. This makes the 1+1 scheme less efficient in bandwidth utilization but more robust. For M:N scheme, M pre-allocated backup paths are used to support N primary paths. In the M:N scheme, the backup paths do not transport backup data at the same time while the primary paths are in operation—only in case of failure do the backup paths become active. The M:N scheme has two special cases: 1:N and 1:1, in which only one dedicated backup path is used to support N or 1 primary path, respectively. The 1+1 scheme is more common for TDM networks. A side effect of M:N protection scheme is that other types of traffic, which do not require high availability, can be placed on the backup link and preempted when needed for the application that was assigned to the protection link.

For restoration, two approaches are proposed in the industry: local repair and global repair. Local repair is applied to a single LSP at path level when a link or node failed. The problem with a link is worked around by *patching* the path in the area where the fault has occurred. In this case, the adjacent node(s) takes action to "detour" the LSP around the failure without the end-to-end re-route. Global repair, on the other hand, provides for a brand new end-to-end backup path. Local repairs are more appropriate for larger networks where fault notification propagation can slow down the recovery time.

Providing High Availability in GMPLS

A key attribute of GMPLS suite of protocols is the ability to enable automated fault-management in network operation. A fault in one type of the network must be isolated and resolved separately from other networks. This is a very important feature for end-to-end LSPs that are tunneled in other LSPs that require higher degrees of reliability along the hierarchy.

In general, four distinct actions are taken to completely resolve a fault:

1. *Fault Detection*, which has to be handled by the layer that is closest to the fault, for example, loss of buffer space for reception of data at the IP layer.
2. *Fault Localization*, which determines exactly where the failure has occurred. The LMP, which runs between two neighboring nodes, includes procedures to localize link-level failures. This is possible because GMPLS allows physical separation of the data bearing channels and control (signaling/routing/management) channels. Specifically, LMP uses the "ChannelFail" message for this purpose, which can be used in pure optical or mixed opto-electrical networks.
3. *Fault Notification*, which takes place once the fault has been localized. It will be up to the node to determine whether a fault requires path recovery. Fault notification is achieved by sending "Notify" messages in RSVP–traffic engineering (RSVP–TE) protocol extensions. This message is sent to the head-end label-switched router (LSR) to perform a global repair or an LSR along the LSP to carry out a local repair.
4. *Fault Resolution*, which takes place either by a switchover to a healthy pre-calculated, pre-assigned protection path or by creation of a new path. This is covered in more detail next.

GMPLS Protection Scheme

GMPLS provides protection against a failed channel (or link) between two adjacent nodes (*span protection*) as well as end-to-end protection (*path protection*). The OSPF and IS–IS extensions for GMPLS advertise the "link-protection type" parameter to include span protection while the route is being computed. After the route is computed, signaling is carried out via RSVP–TE or CR–LDP to establish the backup paths. In constraint-based shortest path first (CSPF), each traffic-engineered (TE) link includes a protection type that can be configured, e.g., unprotected, 1:1, 1+1. The protection type for a given TE link is advertised along with other TE parameters and also stored in the TE database on each node in the network. When an LSP request is received, it includes the "protection type" parameter, which has one of the values, e.g., 1:1. During the CSPF procedure, the routing module performs a TE database lookup to try to find the matching TE links for the protection type requested.

For span protection, 1+1 or M:N protection schemes are provided by establishing secondary paths through the network and using signaling messages to switch from the failed primary path to the secondary path. This kind of protection attempts to emulate SONET's APS, shown in *Table 1*. Recovery time of 50 ms or less is the target.

For end-to-end path protections, the primary and secondary paths are computed and signaled to indicate that the two paths share reservations. End-to-end protection requires that the notification signals travel to the egress node, which will then provide the switchover to the backup path. Since the propagation delay may prove to be too large for some applications, a more localized repair scheme must be utilized to bypass a failure along the path.

Shared risk link group is an optional mechanism that allows establishing backup LSPs that have none or minimal links in common with the primary LSP. This is achieved in the routing extension of OSPF/IS–IS.

GMPLS Restoration Scheme

As mentioned, *restoration* of a failed path refers to dynamic establishment of a backup path. This process requires dynamic allocation of resources, which in turn requires route calculation and signaling for the new path. Two different restoration methods are given: line and path. Line restoration finds an alternate route at an intermediate node. Path restoration is initiated at the source node to route around a failed path anywhere within the path for the specific LSP. As mentioned earlier, they are also referred to as local and global repair, respectively. In addition, backup paths can be pre-calculated or computed after a fault has occurred. The choice depends on the level of QoS that is required by the application. *Figure 2*[1] provides a snapshot of the possible protection mechanisms that could be used in a mixed network.

To expedite GMPLS restoration, recent Internet Engineering Task Force (IETF) drafts provide the notions of

FIGURE 2

Protection and Restoration in GMPLS

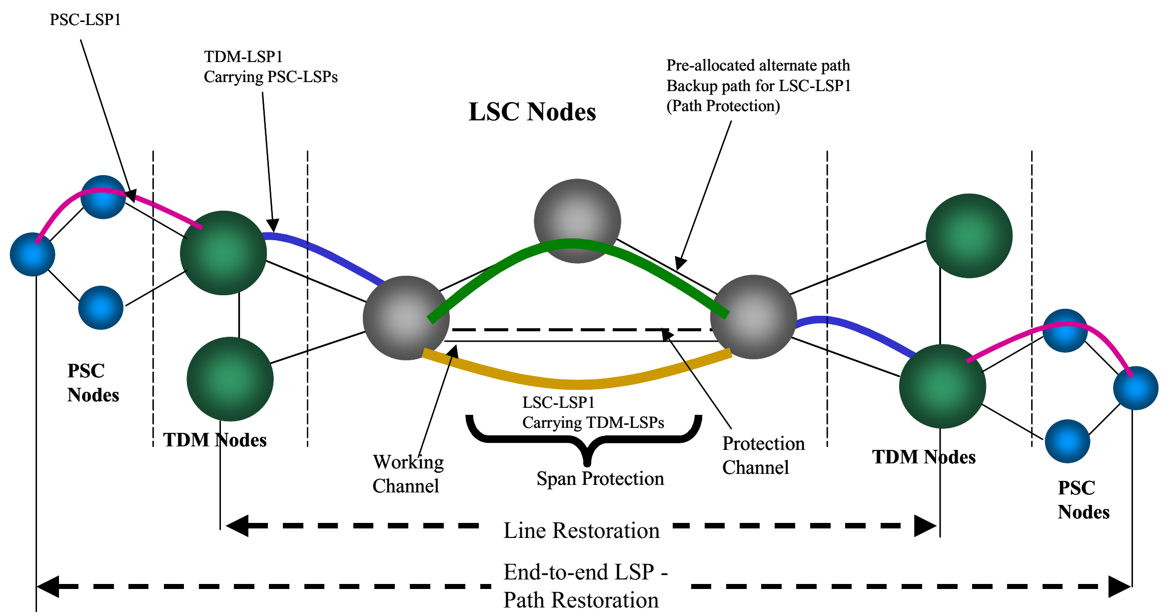

"fast reroute," "detours," and "bypass" [6, 8] in which backup routes are pre-calculated and detours are used along the path of an LSP to remedy a fault by extending the RSVP–TE signaling protocol. The use of a "detour" parameter allows the nodes to differentiate the main LSP from the *detour* LSP. Reference [8] provides a method for unsignaled backup paths.

Both, "detour" and "bypass" LSPs require that the ingress LSR include the "FAST_REROUTE" parameter and use two flags that indicate the need for local protection in the SESSION_ATTRIBUTE object in the RSVP–TE PATH message. This allows the nodes on the primary LSP, which support either "detour" or "bypass," to know that the primary path requires protection.

Currently, there are several proprietary GMPLS LSP protection/restoration/fast reroute drafts that have been recommended to the IETF. The problem is that none of them cover the computation of maximally disjoint backup paths. There are, however, many methods to achieve this. Reference [12] provides a generic method for this calculation, in which the main objective of the given algorithm is to exploit sharing of backup LSPs to reduce the total amount of bandwidth that is consumed by the LSPs.

It is important to note that while the backup path is calculated, there is no guarantee that it will be available once needed, as it could have had failures itself. This could be detrimental to devices that may require a significant order of time to set up their internal switching mechanism, such as a DCS. To remedy this shortcoming, one solution could be having the CSPF to periodically compute the backup path at the egress node to make sure of its availability to reduce restoration time.

Conclusion

The suite of protocols that comprise GMPLS promise to bring true network automation to a variety of dissimilar networks, electrical and optical. Protection schemes are an integral part of GMPLS operations. These schemes use the widely available SONET protection methods as a benchmark but go beyond that to provide a more granular protection level for those applications that do not require very high availability levels. GMPLS' capability to provide SONET–like protection as well as restoration, via dynamic backup path establishment, allows for highly efficient utilization of network resources. Service providers will be able to provide various levels of QoS to their customers and consequently increase their revenue base.

References

[1] P. Ashwood-Smith et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," Internet draft, work in progress (June 2001), draft-many-gmpls-architecture-00.txt.

[2] J. Lang et al., "Link Management Protocol," Internet draft, work in progress (March 2001), draft-lang-mpls-lmp-02.txt.

[3] P. Ashwood-Smith et al., "Generalized MPLS Signaling – RSVP–TE Extensions," Internet draft, work in progress (July 2001), draft-ietf-mpls-generalized-rsvp-te-04.txt.

[4] W.S. Lai et al., "Network Hierarchy and Multilayer Survivability," Internet draft, work in progress (October 2001), draft-ietf-tewg-restore-hierarchy-00.txt.

[5] V. Sharma et al., "Framework for MPLS–Based Recovery," Internet draft, work in progress (July 2001), draft-ietf-mpls-recovery-frmwrk-03.txt.

[6] D.H. Gan et al., "A Method for MPLS LSP Fast-Reroute Using RSVP Detours," Internet draft, work in progress (July 2001), draft-gan-fast-reroute-00.txt.

[7] K. Kompella et al., "OSPF Extensions in Support of Generalized MPLS," Internet draft, work in progress (July 2001), draft-kompella-ospf-gmpls-extensions-02.txt.

[8] A. Atlas et al., "MPLS RSVP–TE Interoperability for Local Protection/Fast Reroute," Internet draft, work in progress (July 2001), draft-atlas-rsvp-local-protect-interop-01.txt.

[9] A. Banerjee et al., "Generalized Multiprotocol Label Switching: An Overview of Signaling Enhancements, and Recovery Techniques," *IEEE Communication Magazine,* Vol. 39, No. 7 (July 2001).

[10] D. Awduche and Y. Rekhter, "Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects," *IEEE Communication Magazine,* Vol. 39, No. 3 (March 2001).

[11] A. Banerjee et al., "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements," *IEEE Communication Magazine,* Vol. 39, No. 1 (January 2001).

[12] M Kodialam and T. V. Lakshman, "Dynamic routing of bandwidth guaranteed tunnels with restoration," Proceedings of the Conference on Computer Communications (IEEE Infocom) (March 2000).

Notes

1. In *Figure 2*, PSC is defined as *packet-switched capable* and LSC is defined as *lambda-switched capable*.

Service Integration: Marrying SLAs to Service-Provider Business Applications

Joel Halpern
Co-Founder and Chief Technology Officer
Longitude Systems

The Service-Provider Crisis

Service providers are under continuous pressure to offer the most competitive services and to attract and retain customers. Both, however, cost considerable time and money. Adding to service providers' bottom-line issues are infrastructure costs that are rising faster than revenue, an inability to differentiate service offerings, marginally integrated operational environments, and service-introduction delays. Service-level agreements (SLA) are being proposed as a method of increasing revenues. Unless they are resolved, these problems will increase costs even faster. This paper discusses the issues of information inconsistency in the SLA–focused environment.

The Promise of Win-Win SLAs

One goal for service providers is to create SLAs that benefit both them and their end customers. A win-win SLA offers service providers the following:

- A foundation for new value-added service offerings
- Strong services differentiation
- Significant new revenue opportunities

A win-win SLA offers end customers the following:

- Reliable operation of mission-critical applications
- Higher employee productivity
- The opportunity to introduce new business models

Why the Promise Is Unfulfilled

Every service provider wants to attain these goals, but many are not succeeding. They have the tools needed to thrive, but they lack the appropriate mixture and use of these tools. Point solutions are very important; without them a provider cannot deliver the SLAs. But it is not enough. A service provider must tie together SLA monitoring, SLA delivery techniques, ordering techniques, billing, and provisioning. An enterprise-oriented solution is needed to work with all these aspects. Without that, providers have true operational inefficiency rather than true operational efficiency.

For instance, the following example of swivel-chair management inside a major Internet service provider (ISP) describes one inefficient operational environment. The provider has contracts with customers and is trying to make sure that billing is enabled at the right time and that the provisioning of routers is done at the right time. Then the provider adds proactive service monitoring, service enforcement, and service maintenance mechanisms. The provider must notify all these systems of the right things at the right times. This would require one employee to sit at two or three terminals, trying to cut and paste and reenter the correct data. Perhaps understandably, the employee gets it wrong. If ISP data in the current environment is only 80 percent accurate (and 80 percent is considered a generous estimate for the accuracy of an ISP), imagine what will happen when providers try to roll out these kinds of additional services.

Service providers must deliver SLAs, but they need an environment in which they can deliver them. They must be able to enforce network and business rules, achieve data integrity, and—most important—scale. In other words, to be able to deliver and monitor SLAs service providers must move to an environment that can tie together the service offerings, the network capabilities, and the business and network rules they have in place (see *Figure 1*).

Service Integration

A service provider needs something that can detail the kinds of service offerings it has as well as its network capabilities. The provider then must relate that to what it actually has deployed so that when it rolls out another kind of service, the network understands which of these SLA monitoring and enforcement techniques must be notified so that all the pieces can come together. The system must know

FIGURE 1

Service Integration Model

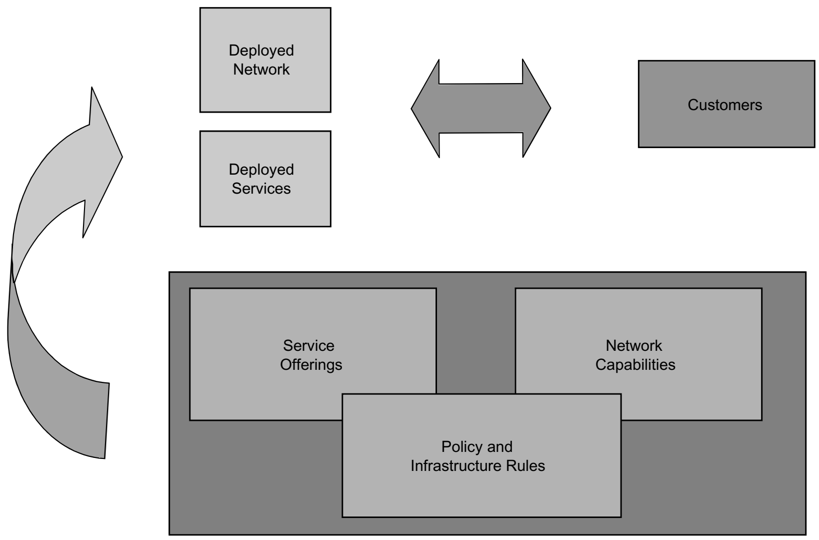

what network infrastructure is available, how to provision the service, when it is correct to notify billing, and so forth. A service provider needs an operational model that is not only a data repository, but that can also tie the operational behavior together so that it becomes a truly intelligent service integration, bringing together the network engineering processes with all other activation and monitoring processes (see *Figure 2*).

However, this does not cover all the functional areas necessary for a service provider to operate successfully. Many other operations support systems (OSS) have evolved (e.g., order entry, billing, and customer care, to name a few) that must be integrated to deliver the kinds of services and SLAs that customers expect.

Total SLA Implementation Solution

Service providers must be able to define SLA–based services independent of their network. They need to deliver those services efficiently and effectively, but they also must integrate them with the rest of their business process. A

FIGURE 2

SLAs in the Operational Environment

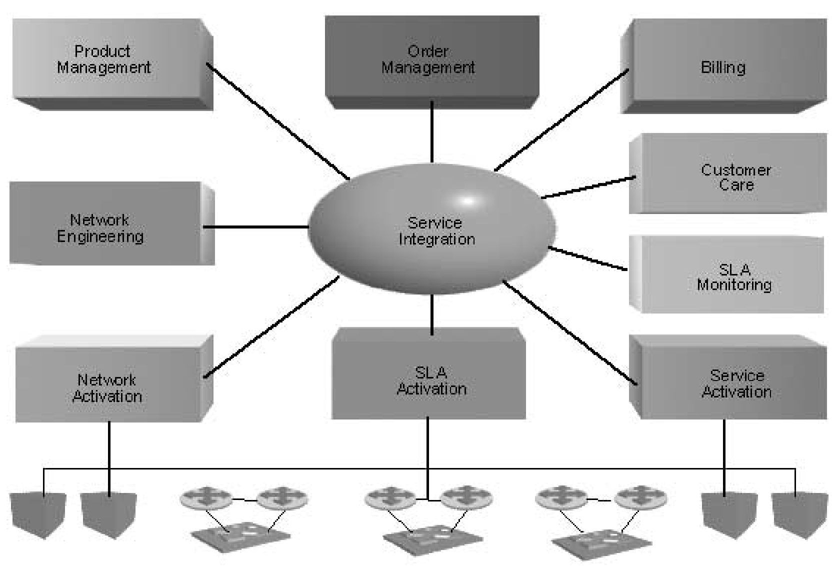

provider's services must tie into billing, customer care, order entry, provisioning, etc.

Services must evolve because the key to success is the ability to innovate and introduce new services continuously. Any service a provider introduces next week—even if it is really innovative—will, in turn, be introduced by a competitor in six months. A provider should be able to introduce the next one after that and to change its business processes either as new components become available or as they offer new services. Successful service providers must evolve continuously, while maintaining data integrity.

For example, when ISPs began buying Juniper routers, many stored the data in the Cisco router table. Why? Because that was the table they already had. This meant they were losing information before they even started because the two are not interchangeable.

In fact, service providers need to take advantage of such differences to gain and maintain competitive advantages. They need a system that lets them greatly improve data integrity, operate efficiently, and scale. No company wants to be a small operation. Any organization in the service-provider industry wants to be a major player because that is how it gets traction. Many providers lament that it is difficult to deliver SLAs when it is over another organization's network. Yet, if a provider has a bigger network, it can control more, and that requires scaling up and growing. The endgame is possessing an operational environment with the intelligence to drive the business processes to achieve efficiency and customer satisfaction.

Call Control in the New Public Network: An Overview of Network Elements and Signaling Protocols

Andrew Hendry
Product-Marketing Manager
Oresis Communications

Introduction

The monolithic circuit switches of the traditional public switched telephone network (PSTN) are well known. There are more than 19,000 Class-5 switches deployed in central offices (COs) and more than 500 Class-4 tandem switches deployed in regional offices across the United States alone. While the annual growth rate of these voice switches has recently been slow and steady, the traditional voice switching architecture will soon undergo an evolution similar to that of the computing architecture in the early 1980s, in which advanced personal computer (PC) workstations running third-party software replaced vertically integrated proprietary mainframes.

Carriers have begun to migrate to next-generation packet-switched voice networks because of intensifying competition, significant cost efficiencies, and new service revenue opportunities. Multiple vendors, both new and incumbent, that specialize in switching hardware, call-control software, and new applications software are beginning to offer a dizzying array of different solutions. Nevertheless, astute choices can be made to optimize the performance of the packet voice solution and the return on investment (ROI) of the purchase decisions.

The Softswitch Architecture

In an environment of intensifying competition for customers among incumbent and competitive carriers, next-generation packet-switched voice networks offer several fundamental benefits: significantly reduced capital and operation costs, less reliance on overlay networks, simplified service provisioning, and new services creating additional revenue and differentiation.

As shown in *Figure 1*, the four key elements in the softswitch architecture are as follows:

- Media gateway
- Softswitch (or media gateway controller)
- Signaling gateway (functionality may be integrated in softswitch)
- Application server

Media Gateway

The media gateway is an access/trunking switch that converts and switches time division multiplexing (TDM) voice traffic and packet voice traffic and provides interworking between the two types of interfaces. The packet voice traffic can be either Internet protocol (IP) or asynchronous transfer mode (ATM), depending on the capabilities of the media gateway and the carrier requirements. The media gateway should be capable of providing interworking between multiple access protocols: for example, any pair of TDM, ATM, and IP.

Important factors for media gateways are as follows:

- *Scalability:* Media gateways must accommodate current traffic as well as future growth. Media gateways currently range from 2 to 20 gigabits per second (Gbps) switching capacity. Port capacity for a Class-4/5 switch typically ranges from 10,000 to 100,000 digital signal (DS)-0s.
- *Reliability:* Media gateways must be highly reliable, usually exceeding five-nines availability (99.999%) to be considered carrier-class. This type of reliability requires superior components, quality manufacturing processes, and rigorous product testing. Full fault-tolerance and redundancy are critical to service availability assurance. NEBS Level 3 Certification is required and OSMINE compliance may be required depending on the carrier.
- *Interworking and Interoperability:* Carriers may use media gateways and softswitches from multiple vendors; as a result, interoperable media gateway control protocol (MGCP)/media gateway control (MEGACO)/session initiation protocol (SIP) protocols are necessary. Bearer

FIGURE 1

The Softswitch Architecture

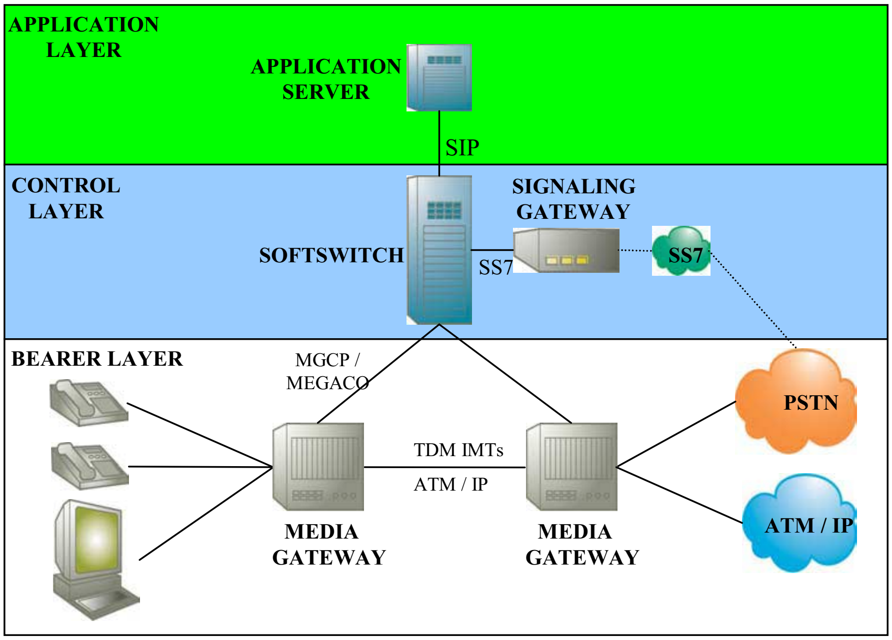

interfaces (TDM, ATM, IP) and signaling protocols (private network-to-network interface [PNNI], bearer independent call control [BICC]) interworking are also essential.

- *Footprint:* The media-gateway footprint drives occupancy costs. More scalable products will accommodate denser modules and typically reduce occupancy costs.
- *Cost:* Carriers deploying media gateways will always consider a cost-per-port comparison. Density is typically an important factor in lowering cost-per-port.

Most media gateways have been developed initially to transport voice using either ATM cells or IP packets. While many industry analysts think that carriers will migrate to IP/multiprotocol label switching (MPLS), for now most carriers rely on ATM to provide the carrier-class quality of service (QoS) required to deliver packet voice services. Either way, a media gateway that meets industry standards is critical.

Softswitch

The softswitch, or media gateway controller, provides the intelligent call control for the voice network. The softswitch also provides another essential function in the context of packet networks: It provides the interworking between the different signaling on the TDM and on the packet portions of the network, e.g., integrated service digital network user part (ISUP) and MGCP protocols. The softswitch may have a Class-4 and/or Class-5 feature set, depending on its role in the network.

Important factors for softswitches are as follows:

- Feature Set: The number and type of Class-4 and/or Class-5 features provided.
- System Capacity: Voice processing capability, typically measured as voice calls per second or busy-hour call attempts (BHCAs), per system.
- Internetworking: Softswitches must support and interwork multiple signaling protocols to provide more flexibility to the carrier deploying its system of choice.
- Reliability: Softswitches must provide carrier-class reliability. Full fault-tolerance and redundancy is critical to service-availability assurance. Fault tolerance between the softswitch and the media gateway is also essential.
- Interoperability: Softswitches must interoperate with a wide range of other network elements, especially other softswitches, because their role is central to the functionality of the packet voice infrastructure. Carriers view interoperability as the top concern in the purchasing decision for new equipment, as seen in *Figure 2*.

The softswitch can be either physically separate from or physically integrated with the media gateway. In the former, a general-purpose server platform such as the Sun Netra or Compaq PCI server provides call control, and, in the latter, an optimized module fits within a standard slot on the media-gateway chassis.

Although most vendors use external server platforms, a few vendors use the optimized softswitch modules. Physically integrated softswitches provide multiple advantages over these general-purpose platforms: simpler interoperability, higher call capacity, smaller footprint, better redundancy, lower cost, and more efficient operations.

Another important feature of certain softswitches is the capability to provide call control for multiple media gateways distributed across the packet voice network. While solutions from many vendors require co-location of a

FIGURE 2

Interoperability Tops Carrier Concerns for Next-Generation Equipment

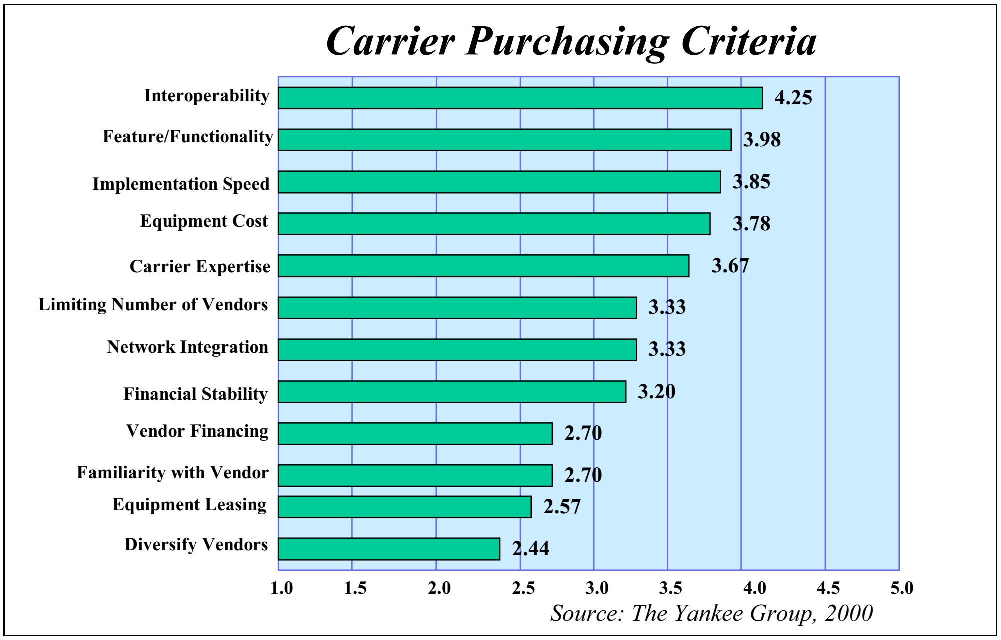

softswitch with every media gateway, a more efficient and cost-effective approach is for remote softswitches to control bearer voice traffic on multiple media gateways located in separate COs.

Signaling Gateway

The signaling-gateway function communicates directly with the signaling transfer point (STP) in an signaling system 7 (SS7) network via SS7 signaling messages in order to provide authorized services and route voice traffic through the bearer traffic network. In some vendor solutions, the signaling gateway is a separate network element but ideally is integrated into the softswitch, which offers lower-cost, simpler operations and a smaller footprint.

Application Server

The application server is a product that provides enhanced services that may include "follow me" roaming, automatic language-translation services, and data-driven voice services. Until now, carriers have relied heavily on the feature set of the traditional circuit switch. In a softswitch environment, however, carriers will have the flexibility to offer new services without replacing or updating the entire packet voice platform.

Alphabet Soup: Signaling Protocols

How the primary network elements in a softswitch architecture communicate with each other is not an easy discussion, because there is a plethora of standards-based and proprietary protocols. Nevertheless, the total number of used protocols is shrinking and can be segmented based on the specific interfaces, as seen in *Table 1*.

PNNI is the standards-based protocol used by media gateways and ATM switches to share routing and signaling information to establish and maintain real-time voice and data connections. The BICC signaling protocol, which is a modified SS7 ISUP protocol, correlated with ATM and PNNI signaling, is used as a signaling interworking function.

TABLE 1

Common Signaling Protocols

Network Elements	Signaling-Protocol Options
Media Gateway–Media Gateway	PNNI, BICC
Media Gateway–Softswitch	MGCP, MEGACO (H.248)
Softswitch–SS7 Network	ISUP, TCAP
Softswitch–Softswitch or Application Server	SIP

MGCP is the signaling protocol between the media gateway and softswitch used until now for carrier deployments. Most vendors' MGCP implementations are proprietary. MEGACO (H.248) is the compromise International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) communications protocol that will eventually replace MGCP. MEGACO is still under development by most vendors and has not been widely tested for interoperability. Consequently, MEGACO is still too new for carrier deployments.

ISUP is a SS7 protocol for signaling the parameters and procedures to set up and tear down circuit-switched voice calls between a softswitch/signaling gateway and an STP. Transactional capabilities applications part (TCAP) is used to exchange control-related messages between a softswitch/signaling gateway and application databases (service control point [SCP]).

SIP is an Internet Engineering Task Force (IETF) protocol for transporting call control, authentication, and other signaling messages among softswitches and other devices through an IP network.

Conclusion

The migration to packet-switched voice networks will be an evolution, not a revolution. The most likely scenario is for carriers to pursue a "cap and replace" strategy for Class-4 tandem switches, followed later this decade by Class-5 switch replacement. With so many emerging network elements and communications protocols, the vision of the new public network can be somewhat daunting. But similar to the revolution from mainframes to PC workstations, significant capital cost and operational savings will prove to spur carriers and service providers to make the migration. With the range of packet voice infrastructure products currently being developed and launched, this segment is definitely one of the most exciting in telecommunications.

APPENDIX

Purchase Recommendations

Media Gateway	*Purchase Recommendation:* Select a media gateway that offers scalable port density, has a high initial non-blocking switching capacity, offers multiple physical interfaces, has proven reliability, has proven interoperability with external softswitches, occupies a small footprint, and begins with ATM while supporting IP/MPLS longer term.
Softswitch/Signaling Gateway	*Purchase Recommendation:* Select a softswitch that can meet call-processing capacity; has a desired feature set; has proven reliability; has proven interoperability with media gateways, other softswitches, and application servers; and incorporates the signaling gateway functionality. A physically integrated softswitch capable of supporting call control with multiple distributed media gateways offers distinct performance and value advantages.
Application Server	*Purchase Recommendation:* Select an application server that provides attractive new services, scales to meet customer needs, and has proven interoperability with other softswitches and application servers.

Operationalizing Voice over Broadband: Integrating the Operations Support System

Razi Imam
Vice President, Marketing and Business Development
CoManage Corporation

This paper focuses on the challenge of ending the fragmentation problem in the United States for voice over digital subscriber line (VoDSL) and rethinking the operations support systems (OSS) architecture.

The Challenges

OSS fragmentation is a significant issue for today's service providers. Providers currently take too long to set up, provision, activate, and offer services. Other considerations include how network convergence will affect the next-generation services and how service providers can leverage and deliver high-margin services on this new infrastructure. Finally, providers must determine how to address customer expectations, decide which management functions need to be integrated, and deliver the quality of service (QoS) required for next-generation voice services.

Network Element Convergence: The Collapsing Infrastructure

In the OSS fragmentation picture today (see *Figure 1*), service providers are dealing with multiple types of networks, such as asynchronous transfer mode (ATM) core networks, Internet protocol (IP) networks, or synchronous optical network dense wavelength division multiplexing networks (SONET DWDM). These networks are made up of multiple boxes connected independently. Each box employs its own proprietary technologies. Each box has its own method of being managed, whether through simple network management protocol (SNMP) transaction language 1 (TL1), common object request broker architecture (CORBA) or proprietary element management system (EMS). From an OSS perspective, this situation creates a nightmare because operations personnel have to deal with different types of networks and equipment when they are setting up services. Using Telnet (a command-driven, text-based interface) to manage these boxes and set up services becomes a very difficult challenge, given the multitude of different equipment types found in a typical carrier's network.

The OSS Fragmentation Problem

Until now, the solution employed by many has been to acquire various software to perform individual network management tasks (see *Figures 2* and *3*). Order management software is purchased to create a link between the customer- and network-facing tasks. Providers attempt to reduce activation bottlenecks with provisioning and activation software that physically talks to the equipment and creates an end-to-end service. They buy provisioning and fault-monitoring software to know if faults occur and to create trouble tickets. Performance-monitoring software is especially necessary if the provider is selling guaranteed service-level agreements (SLA). In all, providers are purchasing a multitude of costly software applications just for network and service management. Many carriers also invest in other customer-facing and operations applications, such as billing, inventory management, interconnection, workforce management, and trouble ticketing.

Every piece of equipment in the network has its own way of being managed. Software providers all publish unique applications programming interfaces (API), if at all. Database schemas are not only different, but often obscure. And no standards have ever been agreed upon to unify data within OSS software.

Many software applications are required to activate a single service. On Canadian service provider studied can take 35 to 40 days to activate a transparent local-area network (TLAN) service because it is necessary to Telnet into 30 to 40 different boxes to set up the service. With such complexity, many services are error-prone and never get completely activated in about 50 percent of activation attempts. The reality is that the activation process is extremely fragmented.

Real problems occur when the above reality is coupled with customer expectations. Customers want fast service turn-up. They want accurate billing. Other customer expectations include service quality, application-driven QoS, greater

FIGURE 1

Network Element Convergence—The Collapsing Infrastructure

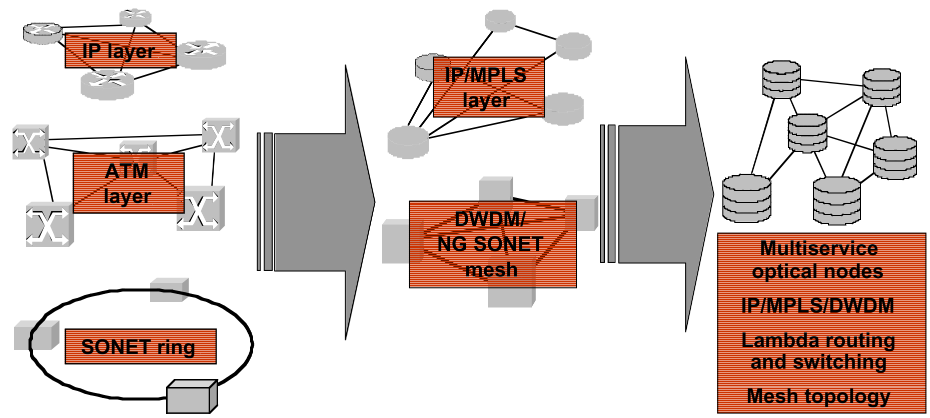

service-provider responsiveness, and proactive customer relations in the event of outages. Customers evaluate service providers in these terms. Meeting customer expectations is the impetus for customer-aware service management.

The Converging OSS

One solution is the convergence of OSS (see *Figure 4*). Some companies have begun to develop the next-generation VoDSL services or similar broadband services—but instead of maintaining their legacy systems, these providers are considering software applications that have been designed from the ground-up to help provide these services faster. For example, a single application could give the provider the ability to discover the network, provision the network, and activate the services. The same application that monitors the services also monitors the network, providing a unified view. A unified view is important because it builds a network knowledge base within a single application.

A single application allows providers to be proactive. If a fault occurs on a certain port, it is possible to determine what services and customers are affected by the fault. The provider can connect to the trouble-ticketing packages and the customer-billing packages to report that the fault has occurred. Similarly, if a new service is activated, the application can direct billing applications to start charging for the new services.

Customer-Aware Service Management

Customer-aware service management has four requirements:

1. Awareness and management of the customer's equipment, including all edge and core gear.
2. Awareness and management of the customer's services. The provider must have end-to-end views and control of individual connections and overall QoS.
3. Awareness and management of the customer's problems to permit rapid responses and reduce customer churn.

FIGURE 2

Converging Management Systems

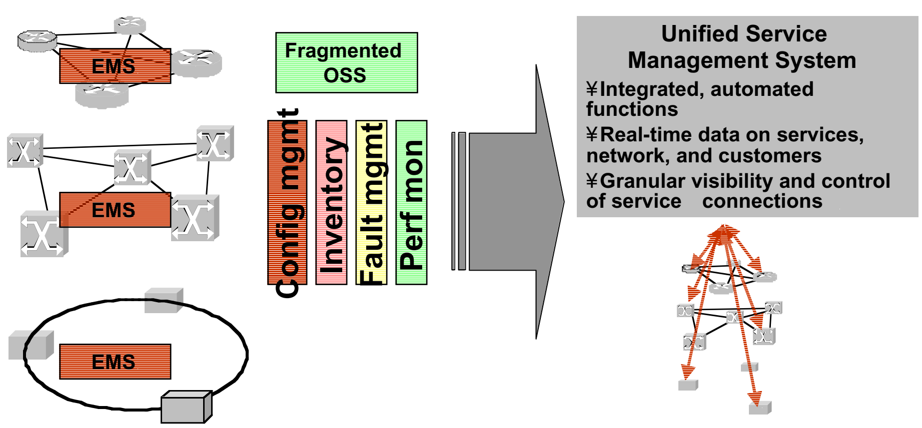

FIGURE 3

The OSS Fragmentation Problem

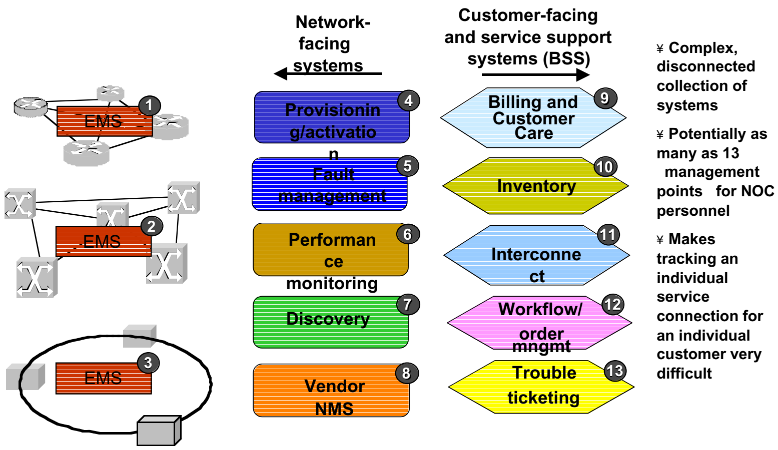

4. Awareness and management of the customer's service quality to offer meaningful SLAs and required QoS.

Building the Customer-Aware OSS

To address these four requirements and fully leverage the network infrastructure, a network-facing OSS must have several characteristics. It must have integrated functionality. It must have tight coupling between QoS, SLAs, and individual service-control functions. Because the DSL infrastructure is composed of a wide variety of equipment from multiple vendors, it must provide multivendor support.

Creating a unified OSS architecture is complex. Hardware vendor solutions are robust but typically only manage their own equipment. Software solutions from third-party vendors tend to be generic and only address one aspect of network or service management. The new goal for service providers should be to buy the best-of-breed hardware to set up the network and then buy a single application to manage the network.

Integrated Functionality

A research study done two years ago with approximately 40 different service providers asked them to identify the functions that could drive operations costs down, the functions that could help them reach profitability faster, and the preferred method of delivery for these functions.

First, the providers wanted to integrate. They wanted a single application, and one of its principal components should be discovery. Providers should be able to discover not only their boxes, but also the ports and the cards and how are they interconnected. The providers wanted to be able to see a topological view of the entire network. In addition, because these are next-generation boxes and services, the application should discover all the services that are currently running on

FIGURE 4

The Converging OSS

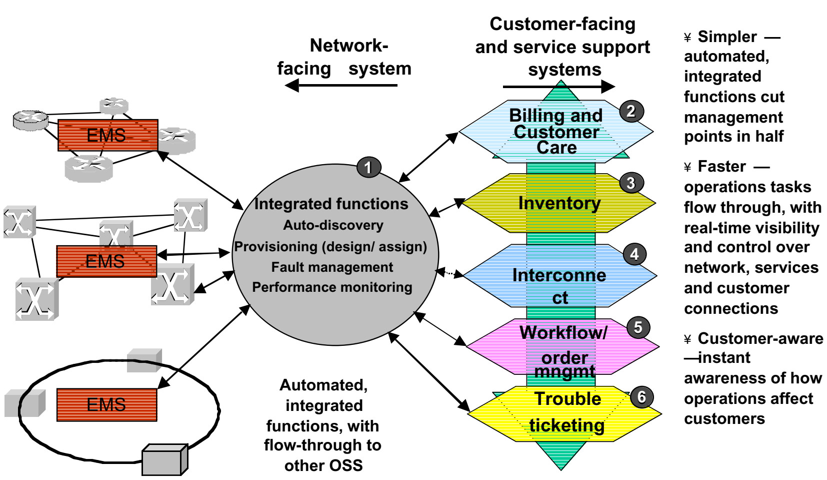

the network so the providers know exactly how much of the total network capacity being used, and what resources are available to activate new services. Discovery was identified as the principal means of reducing operations costs because of the information it can provide. Providers did not want to buy new boxes if they are not required—information that only a robust, real-time discovery system can provide.

Providers also wanted a simple provisioning function that would allow them to select different profiles and offer different levels of service and SLAs. They wanted to be able to click icons to provision and activate those services.

The final components identified by carriers were fault management and performance management. Carriers wanted to know when a fault occurs. They did not want to learn about the fault at a later time, but wanted to be informed by the application when a fault occurs in the network. Ideally, the OSS should be able to identify not only faults in the physical network, but also in the logical or service layer of the network, which is critical to proactive customer service.

In performance management, service providers wanted the ability to understand how services are performing for individual customer at any given moment. This is important for making sure customers' services are meeting SLAs and enabling upsell possibilities.

The key to function integration is knowledge integration. Knowledge integration requires a unified database structure

FIGURE 5

Knowledge Integration Example

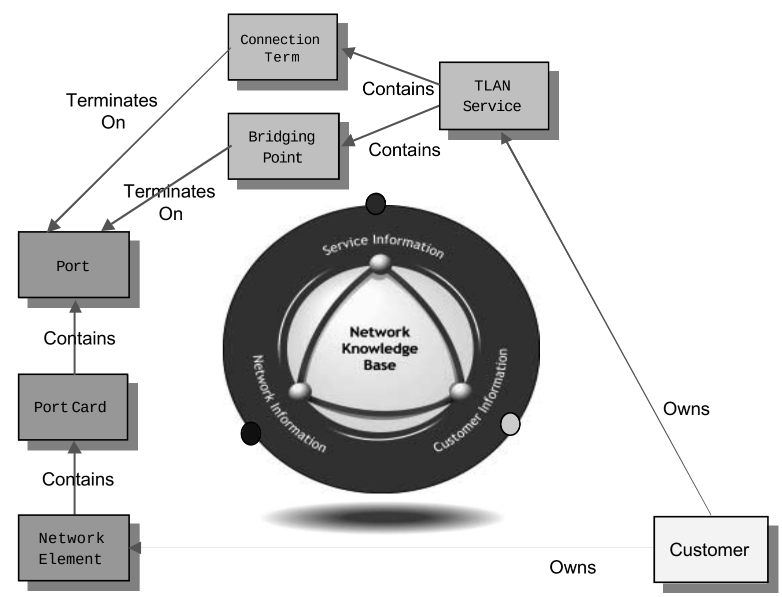

FIGURE 6

Benefit: Single Service/Customer Management for Complex Services

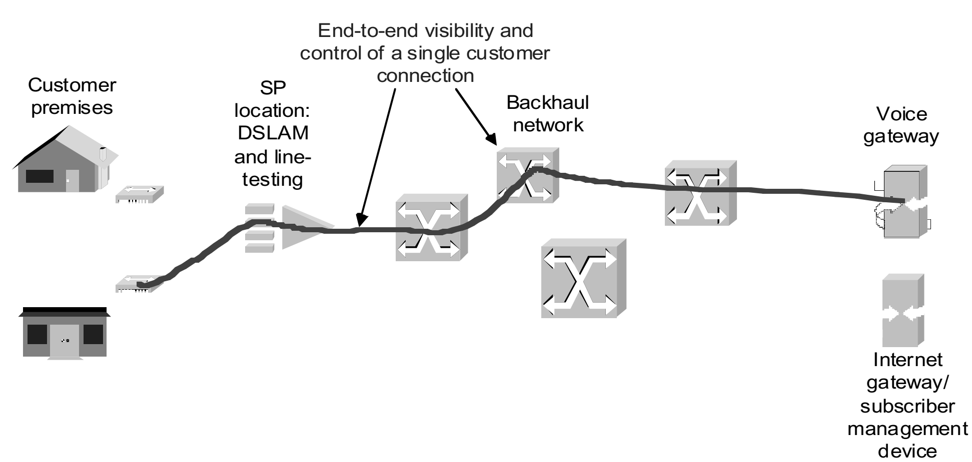

with detailed information on the network, the service, and the customer.

An Example of Knowledge Integration

These four fundamental applications were a prerequisite for offering next-generation data services. Developing this application led to a unified network knowledge base, a database characterized by a correlation that was never possible in the past. *Figure 5* shows all network cards and ports correlated to all services and all customers using those services and ports. This kind of correlation can only happen via a single application with a unified database. If a port goes down, the providers not only know which box it is so they can roll the truck and fix that port, but they also know exactly where it is because of the discovery capability. They know which ports are not functioning, what services are running on those ports, and which customers have been affected.

Such capability gives service providers a competitive advantage because they can call their customers to acknowledge that services are down but are being repaired. This kind of proactive customer service is unheard of in the current market.

Benefit: Single Service/Customer Management for Complex Services

The bold connecting line in *Figure 6* represents all the service elements that comprise a service topology. This figure, an end-to-end example of voice-over-broadband (VoBB) or VoDSL, also shows the hardware components that build the network topology and the customers in the network knowledge base. A single application provides a view from the integrated access device (IAD) to the voice gateway. Providers can provision all of the elements throughout the entire network—the DSLAM, the backhaul network, the voice gateway, and the IAD. Providers can monitor faults on these elements and manage performance over the entire end-to-end service because management is provided by a single application.

Key to Integration Success: Automation

Automation is the key to integration success. Automated discovery provides real-time awareness of network resources, provisioned services, and customers. Automated provisioning has a one-touch service design with simplified QoS assignment and automatic design, activation, and flow-through. Automated fault management starts immediately upon service turn-up, triggering predetermined protection and fail-over processes and speeding root-cause analysis. Automated performance monitoring instantly flags performance issues related to QoS requirements and SLAs, permitting proactive resource allocation to meet traffic demands.

Flow-Through Service Provisioning: QoS/SLA Service Profiles

Providers must be able to easily offer any variety of services to gain a customer's business. The single software application on a next-generation network would allow the provider to pick a service profile that would satisfy the customer or to change the service profile when necessary (*Figure 7*).

For example, a customer may want a gold service profile, which is point-to-point DSL, but wants to upgrade the max-

FIGURE 7

Example of Service Profile

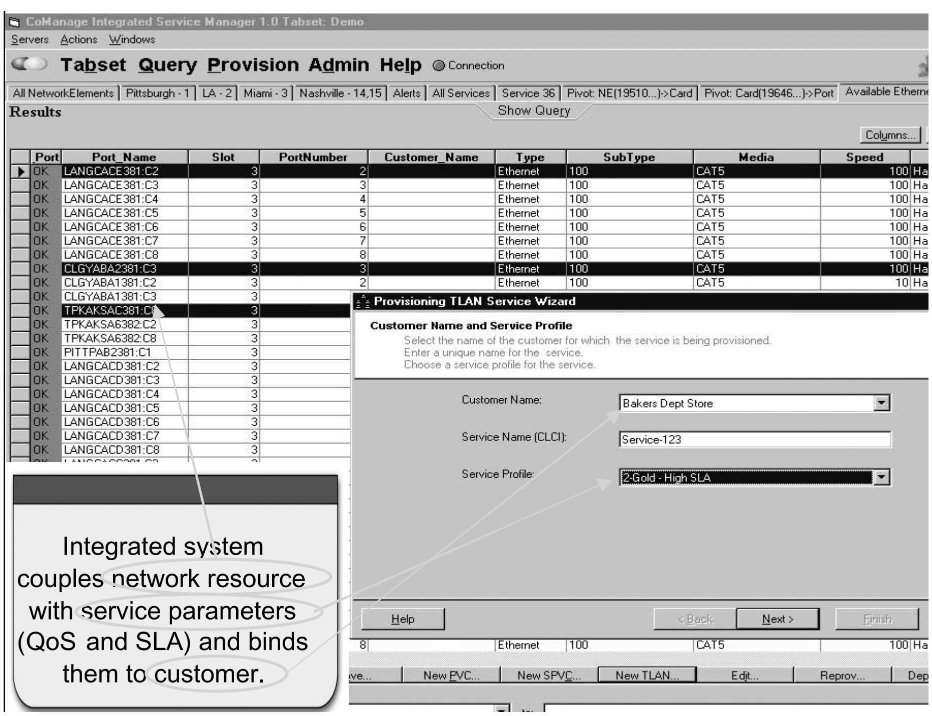

FIGURE 8

Example of Service Topology

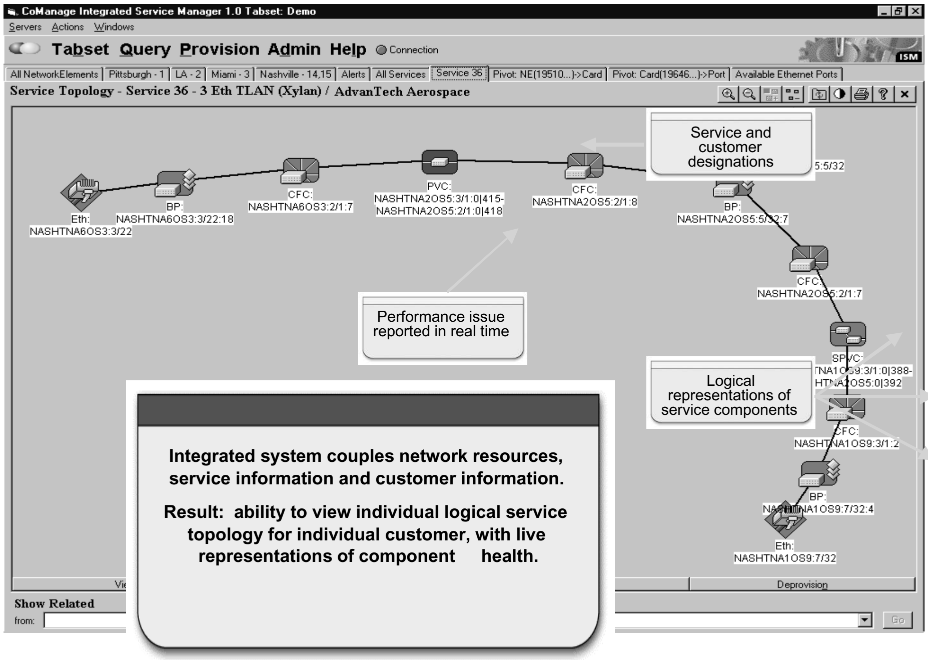

imum bandwidth to 1.2 megabits per second (Mbps). The provider can easily accommodate the request because the software permits real-time modification of these profiles to customize what the sales and marketing staff may sell. Providers then have the flexibility and capability to tell the sales force that although they could only sell one type of service in the past because it took a considerable amount of time to set up the service, the new integrated software solution allows them to sell whatever services can gain the customer's business.

Figure 8 shows a service-topology view of the services. Once the network knowledge base is in place, the provider is able to discover, provision, and fault-monitor the network and performance-manage services.

From a service standpoint, these are next-generation services. The benefit of next-generation service management is

FIGURE 9

Multivendor Support

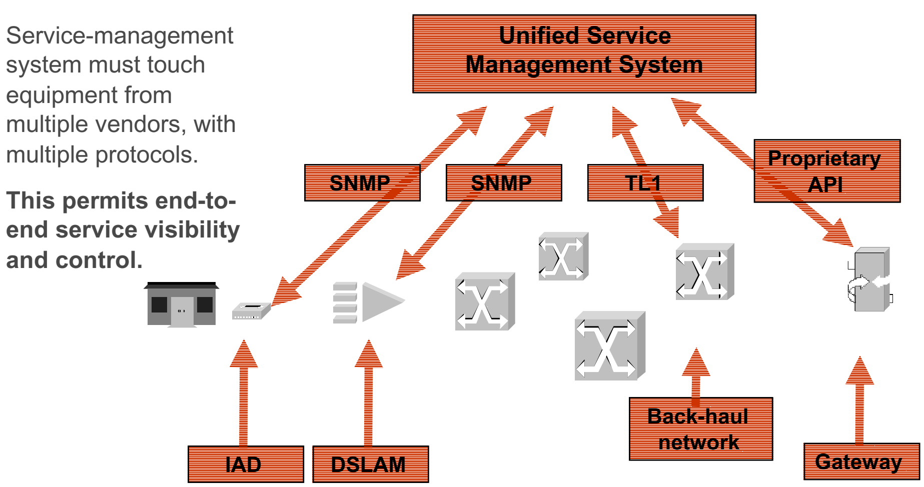

the ability to monitor SLAs and faults and immediately identify customers that are adversely affected quickly and easily. If a customer has a VoDSL service or a transparent LAN service, the provider can see what part of the logical layer has broken down and then fix it or route around it.

Anyone trying to integrate five different vertical applications—especially the network-facing applications such as discovery, provisioning, fault management, and performance—will find that the integration process never stops. It never stops because applications change and versions change. Integration is a continual work in progress. Network-facing applications should be unified into a single application because it provides the maximum benefit of building the network knowledge base, without ongoing systems integration expenditures.

Multivendor Support

Multivendor support allows providers to choose what equipment suits their budgets and needs and to be able to build a network through a multivendor setup (see *Figure 9*). Providers can select best-of-breed DSLAMs and voice gateways without concern for operation support interoperability.

Conclusion

To take advantage of next-generation services, providers need to simplify. Networking-equipment innovations are collapsing the network infrastructure, so service management systems must simplify as well—especially for managing complex DSL infrastructures. The integration of key functions must take place to end OSS fragmentation in these complex environments. Solutions are available that are specifically designed to get providers running on these next-generation services quickly.

Moreover, providers need to enhance the customer experience by giving customers what they want. Service provisioning and assurance systems must be able to associate strong QoS guarantees with specific SLAs and specific customer needs, particularly for toll-quality voice.

Finally, customer awareness is the key to unlocking the power of the next-generation network. To be a more customer-centric service provider, the network-facing applications should contain customer information. The information does not need to be exhaustive, but it should have a correlation of the customer information embedded right at the core of operations. That information gives a provider a tactical, competitive advantage.

Service Assurance: Empowering Operational Excellence

Norman Kincl
Chief Architect, Integrated Service Management
Hewlett-Packard

Introduction: What is Service Assurance?

A service provider will have many different business processes that together define their business. Some of these will be generic processes that any company can be expected to have. Others will be very specific to communications service providers.[1]

Though aspects of this paper may be applicable in other environments, the focus of this paper is service providers. Using the TeleManagement Forum (TMF) definition, service providers are "companies who provide communications and/or data services as a business. Service providers may operate networks, or they may integrate the services of other providers to deliver a total service to their customers."

The TMF has identified three end-to-end business processes that are specific to a service provider:

- Service fulfillment (also called service delivery)
- Service billing (also called service usage)
- Service assurance

In addition, we consider customer-relationship management (CRM) as a separate process. Though not unique to service providers, it is critical.

Figure 1 shows the TMF processes and how they break down into customer-focused activities. Typically, a service provider uses a collection of systems and manual steps to implement these business processes. This collection of systems is often referred to as the operational support system (OSS).

Within the end-to-end process, the TMF defines "service assurance" as the collection, analysis, and processing of network state information correlated with customer assets to ensure that the offered services to the customer have met and exceeded the required quality of service (QoS) and to provide information for recovery action if otherwise.

Two things need to be made clear about service assurance. First, service assurance cannot be considered in isolation of the broader end-to-end process. As shown above, it is an integral part of a larger set of processes and systems. Second, there is value in looking at and focusing on service assurance apart from the other processes. Even though no two people will decompose a system into the same set of functionality, the larger systems, such as an OSS, need to be decomposed if we are to understand the problems and requirements and be able to build a solution.

What Makes a Service?

Looking at service assurance, one question that immediately comes to mind is this: Exactly what are the services we are assuring? This should be defined from the end customer's perspective: a service is something (in this case, communication- or data-related) that the customer is willing to purchase. This could be as simple as a point-to-point leased-line connection or as complex as a complete Internet protocol (IP) virtual private network (VPN) together with a set of hosted applications, such as sales-force automation and unified messaging.

Each service could be composed of different service elements (SE). For example, in the VPN example above, the service provided to the customer could be composed of the following SEs:

- IP access network
- IP VPN
- Unified messaging
- Sales-force automation application service

A products and services catalog defines how the SEs come together to form a service offered to customers. This catalog is a component that needs to be shared across all the service provider processes. The services offered and SEs used will depend on the business model of the service provider.

Figure 2 shows an example of how a service could be composed of several SEs. Each SE is implemented by a component. These components could be:

- Network components, such as traditional network elements (NE) and their contained components
- Platforms, including servers and basic networking utilities, such as domain name server (DNS), dynamic host configuration protocol (DHCP), etc.

FIGURE 1

End-to-End Process Breakdown

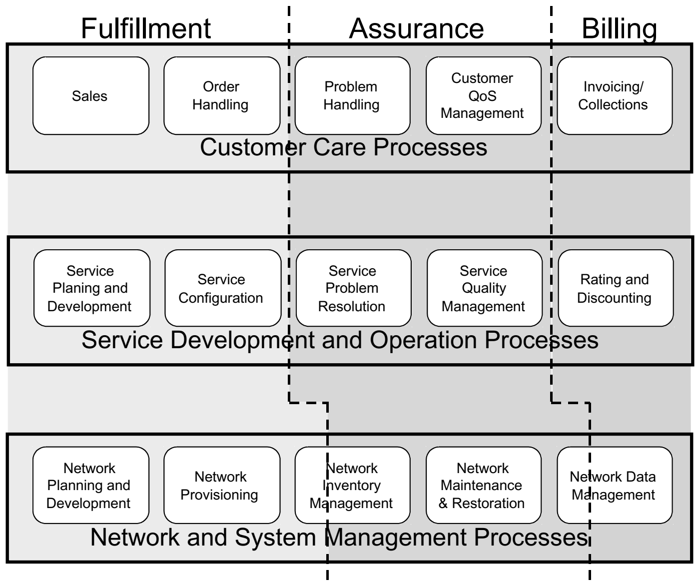

- Applications that provide specific functionality to the customer

The specific categorization of the components is not significant. The important thing here is to understand the potential complexity and variety of the underlying components that implement the SEs.

Merging Models

As the Internet has been growing and becoming an important part of a service provider's business, two cultures have been clashing—the world of the traditional telephony and that of the Internet. This is manifest in many ways, not the least of which is the terminology used in each world. The

FIGURE 2

Example Components of a Service

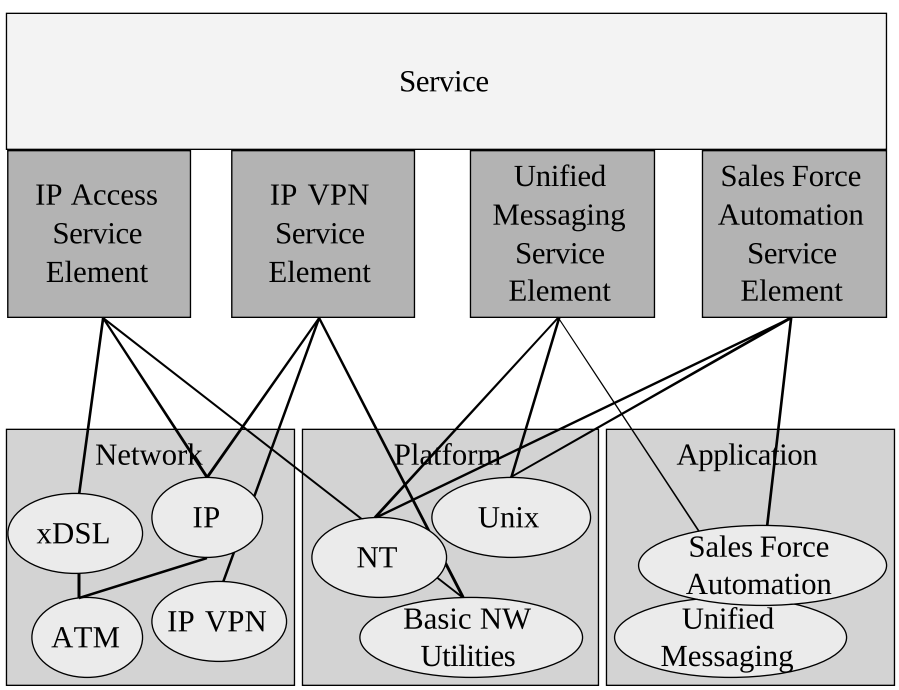

traditional telephony world has a vocabulary that has come out of years of standardization efforts in organizations, such as the International Telecommunication Union (ITU). The Internet world also has a rich heritage, stemming from its academic origins and its "open" approach to development and growth. And just as it began to look like we would have these two groups able to understand each other's vocabulary and communicate, a new complexity arises.

There is now a third set of models, vocabulary, and way of looking at things coming into the fray. As ASPs become part of the service-provider landscape, the culture developed in the information technology (IT) organization also becomes important. In this world, standards, such as the IT infrastructure library (ITIL), define a different vocabulary.

This paper uses concepts and terminology that try to speak to the mixed environment of traditional telephony and the Internet world. The important task of bringing these two cultures together with the IT management culture has not yet been attempted. This should not be interpreted to mean, however, that the concepts outlined in this paper do not apply to service providers focusing on applications or other services that have a traditionally strong IT tie.

Why Service Assurance?

The heavy push throughout all industries and regions into a reliance on electronic commerce (e-commerce) is placing increasing demands upon service providers. It is no longer merely inconvenient if a communications or data service is slow or unavailable. Any disruption in the communications services can have a significant impact on the bottom line of a company.

To address this, companies are increasingly requiring service-level agreements (SLA) from their service providers. Though low cost is always attractive, it is becoming far more critical for a service provider to be able to provide differentiated levels of service.

SLAs

An SLA is a contract between the service provider and its customers, the service subscribers. The SLA defines the service-level aspects of the service contract.

SLAs provide value to both the service provider and the service subscriber. Benefits to the service provider include the following:

- Differentiating service offerings from other providers
- Ensuring that customers have a clear understanding of service expectations
- Providing a clear specification of what workload the providers can expect for the customers
- Defining a common language that can facilitate discussion and resolution of service-related problems
- Faster identification of service problems leading to better service
- Satisfied customers

For the customer, the benefits include the following:

- A guaranteed QoS
- An unambiguous understanding of the guarantees
- A mechanism for auditing the performance of a service provider and holding them accountable when guarantees are not met
- The ability to compare the service quality of various providers and make price-performance tradeoffs
- The opportunity to shift the task of managing and maintaining the service-delivery infrastructure to a provider

Some of the topics an SLA will typically cover are related to the components. These include such areas as availability, throughput, mean time between failure (MTBF), and mean time to repair (MTTR). Since these are directly related to the components supplying the service, they are referred to as component-related SLA parameters.

An SLA can also refer to service-related parameters, such as the time required to provision a new service. Finally, an SLA could also contain customer-related parameters. These could include the amount of time the customer is on hold waiting for a customer-service representative (CSR) or how a customer is notified of service-related problems.

In order to support these SLAs, service providers need to have a system that allows them to monitor and manage the level of service you are providing—a service-assurance system.

What is a Service-Assurance System?

A service-assurance system is a system that allows a service provider to manage the processes and information required for service assurance. They are typically considered part of the OSS. As described in fuller detail later in this paper, a service-assurance system needs to share information with the other systems.

Service assurance takes the view that a service is composed of component pieces. And these component pieces could be provided by the following:

- Some dedicated components, such as a line card or a digital subscriber line (DSL) modem
- Shared components, such as a router or a mail server
- Services outsourced from another service provider, such as a local loop or an IP network

To manage the service, the service provider needs two things. First, they need to manage the individual service components. This typically will come together in a network-management system (NMS). The network view will show the general health of the components (and outsourced services).

Second, the services they provide to their customers need to be managed. The service provider achieves this by putting the status of the component, as discovered by the network view, together with the understanding of how the components work together to provide a service. By combining these, they can view the status of each service provided to their customers. By bringing together customers' expectations as defined in the SLA with the component-, service-, and customer-related SLA parameters, a service provider can manage their customers' satisfaction with their services.

Different Views

A service-assurance system will provide different views of the information to the service provider: namely, network, service, and business views. It is important to understand that, though these views are based on the same underlying data, they each have very different purposes:

- The network view[2] provides a technology-centric view
- The service view looks at things from the perspective of the services offered
- The business view focuses on how well the business is running

In addition, the service-assurance system needs to provide a customer view that provides the service provider's customer with a logical view of the services that they are purchasing.

Network View

The network view provides an understanding of how well all the components that make up services are working. These typically have been provided through an NMS. Currently, these primarily address managing the network components, with some extending into managing platforms. Application management (and much of the work around platform management) has generally been in the domain of IT management. Only recently have these been achieving the scalability or robustness required by service providers.

There are two primary purposes for the network view. First, it indicates problems or trends that could become problems, allowing the service provider to repair them before the customer complains. Second, it provides a strong foundation for building the service view.

Though it is possible, and sometimes beneficial, to use a small number of independent NMSs, the number of systems should be kept to a minimum. Certain activities, such as correlation, become very complex across independent systems.

NMSs have been the focus of many standards. The primary standardization effort has been in the view that the NMSs have of the NEs. There are two primary areas of standardization:

The first area involves the protocols to communicate with the NEs or element managers. SNMP and common management information protocol (CMIP) are the primary standards for communication with network equipment. Other standards, such as transaction language 1 (TL1), as well as proprietary systems are also still widely used.

The second area is the model the network manager has of the NEs. The ITU object model (as defined by the Guidelines for the Definition of Managed Objects [GDMO]) and the SNMP management information base (MIB) are the primary standards. As with protocols, much of the equipment in use today still relies on older models.

Though the telecommunications management network (TMN) model suggests an element-management layer (EML) underneath the network-management layer (NML), this does not always make business sense. Element managers are highly desirable when the NEs are complex, or where the number and diversity of similar NEs provide economic value in creating an abstraction layer. However, when the NEs are very simple or small in number, an element manager may not make sense.

Though a full NMS could include many more areas, we will look at those that fall within the domain of service assurance.

Business Benefits

The service provider can expect a variety of benefits from implementing a service-assurance system that adequately addresses the network view. Among these benefits are the following:

Reducing the cost of operations by doing the following:

- reducing the number of required tasks
- concentrating supervisory activities in a national center during off-hours
- instituting an ability to provide follow-the-sun operations
- maintaining control of the key business processes while allowing outsourcing of others

Maximizing operational efficiency by doing the following:

- allowing operators to monitor and control a larger number of components
- enabling reduction of data coming from the network into only the critical information
- providing network status information to the right people at the right time
- introducing preventive maintenance functions

Maximizing the usage of infrastructure by doing the following:

- monitoring usage capacity
- avoiding congestion through traffic control and load balancing

Managing the capital investment in the infrastructure by doing the following:

- using trend analysis to project future capacity
- understanding utilization of resources

Fault Management

Fault management is the cornerstone of service assurance. It provides the collection and correlation of alarms and other relevant events to provide an accurate view of the health of the network. There are several major functions that need to be addressed by fault management, and they are described in the following sections.

Alarm Collection

Alarm collection is the function of collecting information from the NEs or systems to determine their health. For most modern components, the primary form of alarm collection is to listen for events and map those that correspond to error conditions into alarms.

However, not receiving any error events does not necessarily mean that the NE or system is properly working. For example, a component may suffer a "silent death."[3] There are also some components that may not be able to emit events. Because of this, a robust alarm collection system must be able to poll various devices to ascertain that they are still operating properly.

Not all events will come directly from the NEs or systems. For example, one of the responsibilities of performance management is to provide threshold alarms when specific performance parameters are exceeded. Alarm collection must be able to deal with these as if they were coming directly from the system component.

The final function of alarm collection is to put the alarms into a canonical format. A canonical format allows upstream systems, such as those that provide correlation or service mapping, to process the alarms without needing to understand the details of the component emitting the alarm.

Event Correlation

A complex network can generate a massive amount of alarms. Event correlation reduces this to a manageable level by allowing the operator to focus only on those alarms that represent actual problems. Event correlation should not just reduce the number of alarms; it should also have the capability to increase the information content of the alarms.

The first aspect of event correlation in an NMS should be the ability to eliminate or reduce repeated alarms and transient alarms. This type of correlation should be performed as close to the source of the alarms as possible.

Repeated alarms occur whenever a component emits the same event many times in succession. Two separate processing options should be possible for repeated alarms. First, only the first occurrence is transmitted and all subsequent events are suppressed. Second, the first n occurrences are suppressed, with an alarm being generated only after n+1 events.

Transient alarms compose a set of events occurring within a certain time interval of each other (typically seconds) that indicate a problem and its resolution. Transient alarms are generally completely suppressed. A separate "excessive transients" alarm could be generated if the number of transient alarms exceeds a pre-defined rate.

The second aspect of event correlation is correlating individual alarms into problem conditions. A problem condition is a set of alarms that relate to a single problem. In its most basic form, this will simply present an alarm and the clear for that alarm as a single problem condition. Rather than having two separate alarms displayed, the operator would see one problem condition indicating the current state of the component.

However, problem conditions can become a much more powerful concept with more advanced event correlation tools. For example, as event correlation determines that a set of alarms are all related, with one alarm being the root cause and the rest being related alarms, a problem condition will focus attention on the root-cause problem.

The final aspect of event correlation is its ability to increase the information content of alarms. This is done implicitly in the first two aspects (e.g., an alarm indicating that n+1 events have occurred or a problem condition with a root cause and its related alarms have richer information content than the individual alarms). Robust event correlation should also allow you to explicitly increase the information content. For example, for certain alarms, additional information from a configuration database could augment the event information generated by the component.

Problem Management

Problem management allows operators to act on problems. The minimal function it needs to provide is a display of the problem conditions. This can range from simple text-only displays to comprehensive graphical displays. The critical function is to allow the operators to quickly recognize problems. A highly desirable feature is to provide an indication of the impact of a problem.

The ability to display problem conditions, coupled with the operator's ability to discharge those that have been solved, is often sufficient functionality to address those problems that can be directly resolved by the operations staff. The resolution could be to ignore the problem (e.g., ascertaining that the problem is caused by a thunderstorm affecting the microwave transmission) or to take an action that is of relatively short duration (e.g., sending a reset command to a component).

When the problems involve more people or longer duration, problem management should provide additional support. This support can come in one of two forms. The first, and more traditional option, is integration with a trouble-ticket system. The second option is integration with a process engine.

A trouble-ticket system generally provides support for primarily manual steps. The business processes implemented by a trouble-ticket system are relatively static. Typically, the trouble-ticket system would implement an escalation process. Some trouble-ticket systems will provide assistance to the people solving the problem by allowing them to search for how similar problems were solved.

A process engine is better suited for automating the problem resolution, including automatically performing part of the problem-resolution functions. For example, if a fault indicated the failure of a card, a card-replacement process could be initiated in the process engine. This process could do the following:

1. Enable a backup route, if one exists
2. Verify inventory for the replacement card
3. Order additional cards if inventory is low
4. Ship the card where it is needed
5. Schedule a technician to meet the card
6. Schedule downtime on the equipment
7. Test the replacement card
8. Re-establish the primary route

Process engines possess some distinct benefits, including the following:

- They can allow easy modification of the processes as new insight is gained.
- They can maintain metrics on how the processes are working.
- They can be easily integrated with applications required to solve the problem.
- They can provide for manual steps within a process.

A process engine and a trouble-ticket system are not mutually exclusive. Though there is some overlap in functionality, the overlap is not complete—it may make business sense to have both. For example, the process engine could automatically handle known problems, while a trouble-ticket system could provide the support for problems that have not been understood enough yet to have their resolution automated.

Performance Management

The lack of faults does not necessarily mean that the network or systems are running properly. Though the equipment may be functioning, it may not be performing—the load on it may be such that it is just being asked to do too much. This is where performance management comes in.

Most performance data collected is associated with a "network level" component—an NE, system, or application. In some cases, the performance data could be associated with a service or service component, such as asynchronous transfer mode (ATM) permanent virtual circuit (PVC).

There are two major foci for performance management: real-time (or near–real-time) monitoring and trend analysis. Both are supported by a data-collection function.

Data Collection

Performance-data collection is analogous to alarm collection. Typically, NEs or systems are polled on a regular basis for the performance data. As with alarms, this data needs to be converted to a device-independent format.

It is important to provide a balance in the frequency of the polling for performance data. If the component is polled too frequently, then collecting the performance data will affect the performance. However, not polling frequently enough may not detect developing performance problems until it is too late.

Real-Time Monitoring

The purpose of real-time performance monitoring is to identify performance problems as they occur. For example, the load on a mail server may reach a level where it can no longer service user requests in a timely manner. By comparing current performance values against pre-set thresholds, the performance-management system can determine if there is a problem. When it detects a threshold violation, the performance-management system will raise an alarm and send it, on behalf of the component, to the fault-management system.

In reality, the monitoring is not real time but is subject to the polling interval defined by the data collection. This interval should be in the range of 5 to 15 minutes.

Trend Analysis

The function of trend analysis is to project future system performance based on past performance data. As the data is collected, the trend analysis system analyzes and reduces the data, determining any trends. By projecting the trends outward, the service provider can determine when the system will run out of capacity.

The trend analysis data is often also used in the planning and service-creation portions of the service-delivery processes. Rather than just relying on current trends to project capacity, the service provider can use projections from planned service introductions or system expansions to augment the trend data.

Recovery Management

The primary objective of recovery management is to restore the service components after a failure. Part of this process includes testing and diagnosing the components. To accomplish this, recovery management will involve actions, responses, and pre-planned activities.

These actions and activities could take one of several forms:

- Manual actions are directly selected by the operator and usually used in ad hoc situations or when the complexity of automating a process cannot be justified by its complexity.
- Semi-automatic actions are presented to the operator as suggested actions but are not executed until the operator verifies that the situation warrants the suggested action.
- Automatic actions are automatically executed by the system in response to a pre-defined set of events that has a well-known response associated with it.

Recovery management focuses on restoring failed components. However, it also plays a crucial task in providing service-level restoration. This topic is discussed further in the "Service Problem Resolution" section.

Command Abstraction

A properly designed recovery management system will provide a certain level of abstraction of the components. The commands required to reset or test a component may be complex. It is very likely that these commands will differ between manufacturers, or even releases, of the components.

Abstracting the commands has several major benefits:

- It reduces training requirements.
- It is simpler to introduce new equipment or versions.
- It minimizes operator error.

The third point here should not be taken lightly. A significant percentage[4] of problems occur because of operator error. Situations of operators typing a comma instead of a semicolon and causing a major switch to crash are not unheard of.

It is not the goal of recovery management to fully model every component into an abstract version. Rather, the goal is to implement the "80/20 rule": that is, implement the 20 percent of the commands that are used 80 percent of the time.

Restoration

Depending upon the nature of the fault, restoring the equipment may involve physical changes (e.g., replacing a faulty board) or logical (e.g., issuing a reset command).

When physical tasks are required, a proper trouble-ticketing, workforce-management, or process-flow system needs to track and manage the service restoration process. Though part of the broader restoration process, these systems are not strictly part of recovery management since they play a critical role in so many other processes.

Recovery management deals with the logical parts of the restoration process. Where the restoration tasks can be automated, the value is clear, allowing components to be restored significantly faster than would be possible manually. However, by abstracting the user interface across multiple components, it brings value even if the tasks cannot be (or the service provider does not wish to have them) automated.

Testing

Testing is used to verify the proper operation of a component. Testing may be performed at different times—upon initial commissioning of the component, upon service activation, or during operation to verify its proper function. The first two of these happen during the service-delivery processes and are not discussed in this paper. The later is in the domain of service assurance.

Testing may be initiated for one of several reasons. First, it may be initiated to verify that the component did not suffer a "silent failure," failing without emitting an alarm. As with collecting any management information from the components, a balance needs to be made between querying for management information and allowing the component to perform the task it is supposed to.

Second, other components may indicate that there might be a problem with the component.

Finally, a reported service problem could trigger tests on some or all of the components that provide that service.

Diagnosis

Sometimes a fault indication, such as an alarm, does not sufficiently define where the problem is or what action needs to be taken to restore the component to service. Diagnostics help pinpoint the cause of the fault.

When the diagnostics, but not the recovery, can be automated, it is beneficial for the operator to have the results of the diagnostics "attached" to the problem condition. Sometimes diagnostics can help during the correlation process to help determine if an alarm indicates a new problem, correlates with other alarms, or is not a problem.

Inventory and Configuration Management

Within the service-assurance processes, inventory and configuration management provides a variety of services.

System Details

Often the operators will require details of an NE or system that go beyond those kept by the fault-management system for its purposes. By providing a direct link from the fault-management system to the network-inventory system, the full information kept becomes available. The types of detailed information include the following:

- Decomposition beyond the fault-management containment tree
- Hardware revision levels
- Software revision levels
- Physical location
- Contact information
- Tracking information (e.g., a serial number)

Topology Information

To properly be able to correlate and display fault information, the fault-management process needs to understand the topology of the network and system components. Ideally this information should come directly from the network-inventory system. This is especially important for those components that do not support the concept of auto-discovery.

Auto-Discovery Reconciliation

Some component technologies lend themselves to auto-discovery—it is possible to automatically discover the current existence of resources and their configuration and topology. For these types of resources, the network-inventory system can provide a way to reconcile between the way the network is and the way it was designed to be. This can help detect incorrect configurations, whether done accidentally or maliciously.

Service Mapping

To provide for the service view, a clear mapping between the services and the components that implement those services must exist. This mapping is provided by the network-inventory system.

Spares Inventory

If a failure requires replacement of a hardware component, the recovery process will need to know the availability of spares (with compatible revision levels). This information should come from the network-inventory system.

Traffic and Load Management

The service assurance processes are responsible for the dynamic aspects of providing the services. Beyond recovering from faults (see the "Recovery Management" section in this paper), this dynamic management may include modifying the configuration to improve throughput. With the communications service components, this is generally referred to as traffic management or traffic engineering. With service components that comprise applications or platforms, this is usually referred to as load management or load balancing.

The service-delivery processes are also responsible for a level of traffic and load management. However, during service delivery, static choices are made to determine the best resources to allocate to the service. During service assurance, these choices are modified, within parameter limits set during the delivery processes, to dynamically optimize throughput.

Traffic or load management closely links real-time performance analysis with necessary controls that maximize throughput. The performance parameters measured may be those collected by the performance-management system. However, sometimes the traffic- and load-management systems need to augment the performance-management data collection to meet their specific requirements.

In some cases, such as with an ATM network, this level of control is provided through the NEs and element managers. Even in these situations, it may be desirable to have a higher-level traffic-management system. Such a system would augment the functionality of the lower-level traffic manager in two ways:

- It would provide an independent view into the element manager to provide a view into the decisions being made by the lower-level system.
- More importantly, it would provide the ability to exercise network-level controls that allow traffic and load to be balanced across a broader set of resources.

Service View

The service view helps service providers understand how well they are serving their customers. It looks at the services in total rather than at the components that makes up the service.

In the past few years, many things have been put forth as providing a service view. Though many of these do address some of the aspects required for a service view and do add value, most are ad hoc solutions rather than strong foundations for a new way of doing business.

Probably the most important thing that should be present in a solution implementing the service view is an automatic linkage with the service-delivery processes. Requiring the service-assurance system to be separately configured to monitor the service provided to customers is almost a guarantee that the information will be out of date. Any specific service information required by service assurance should automatically come from the service-delivery processes.

Business Benefits

The benefits provided by the network view primarily focus around cost reduction and operational efficiency. In contrast, the service-view benefits have a more direct impact on the service provider's ability to server its customers. Some of the benefits that the service provider can expect from a service-assurance system that implements a service view include the following:

Improved customer satisfaction and reduction in churn, through an ability to do the following:

- communicate timely and accurate information to customers about their service status
- consistently follow a customer's problems throughout the organization
- prioritize actions based upon the number or importance of customers affected

An ability to offer to customers additional service options, such as the following:

- differentiated service levels
- guaranteed service levels

Service Monitoring

The most basic service-level functionality is service monitoring. Service monitoring takes the network fault and performance information and presents it by service and customer. For example, rather than being told that a particular line card is down, a service-monitoring system could inform the service provider that circuit SJCPA023 for ACME Co. is down.

Instead of equipment alarms and problem conditions provided by the network view, service monitoring provides service alarms and service conditions.

The service-monitoring system needs to be notified of new or modified problem conditions. Whenever the problem condition indicates that it is service affecting, then the service-monitoring system determines which services are affected. It does this by looking at the inventory- and configuration-management information. From this, it determines the affected services and creates service alarms against them. These service alarms correlate into the service condition.

The reverse process should also be possible. When a customer calls a service representative reporting a problem, there is often no way to correlate the reported problem with possible equipment malfunctions. The service-monitoring system should also use the inventory- and configuration-management information to decompose a service into the network components that provide it. The NMS can use this list to test for proper operation.

Two things are critical for service monitoring to succeed. The first is being properly notified by the NMS of the service-affecting faults. The second is an accurate mapping between the service and the components providing the service. The easiest way to maintain this information is by accurately and automatically recording it as the service is created or modified. This is why it is crucial that these service-assurance processes be closely linked with the service-delivery processes.

Service Problem Resolution

Service problem resolution goes beyond simple service monitoring in that it allows the service provider to react to and resolve service problems. This involves two different types of activities.

Resolving service problems consists of fixing the problem with the service. If a failure of a component affects a service, usually the easiest way to restore that service is by fixing that component, a topic discussed earlier in the "Recovery

Management" section. However, if the MTTR for the component will be long or will cause the SLA to be violated, the service provider may need to employ other service-recovery methods. Typically, this will involve activating backup resources or provisioning alternate routes. This is discussed further in the "Service Restoration" section.

Reacting to service-level problems consists of the ability of the service provider to do more than just fix what is wrong. It may include having operations notify customer care of the problems, updating status displays for customers, etc. Generally, this means executing on a pre-defined business process that defines exactly how the service provider wants to react to service problems.

SLA Monitoring

A well-defined SLA needs to be measurable. It needs to be clear to both the service provider and the customer what constitutes acceptable service. This is defined through a set of service-level objectives (SLO). An SLO is a combination of one or more component measurements to which constraints are applied. In addition, the SLO will have an operating period over which it is valid. For example, a set of SLOs defining the performance aspects of an SLA could look as follows:

- service response time < 85 ms between 8:00am and 5:00pm, Monday–Friday
  AND
- service availability > 99.95 percent between the hours of 8:00am and 8:00pm
  AND
- overall compliance > 97 percent over a period of a calendar month
  VALID FOR
- transaction workload < 100 transaction/sec

SLA monitoring is monitoring the service for compliance with these objectives. Rather than just raising an alarm because the response time exceeded 85 ms, an SLA–monitoring system may not raise an alarm when it notices that it is late at night and the response time is allowed to deteriorate. Rather than just reacting to problems (as a simple fault manager might force you to do), or reacting to service problems (as service monitoring might have you do), SLA monitoring allows the reactions to be based on contracts and expectations set with the customer. This allows the service provider to offer his customers differentiated service levels. Instead of having to react the same way in all situations, the service provider reacts appropriately to the service level the customer is paying for.

The SLA–monitoring system could be set up to react before the SLA is violated. By defining a threshold value, the SLA–monitoring system could provide an alarm before the SLA is violated. This would allow the service provider to work proactively to avoid a potentially costly SLA violation.

It is important to note that SLAs define service objectives that go beyond jus the performance or reliability of the service. Some of the objectives will pertain to areas beyond the service components. For example, an SLO may state that calls to customer service will be answered by a CSR in < 10 seconds between 8:00am and 8:00pm, Monday through Friday. Similar SLOs could be created regarding time required to provision a new service or change service parameters.

SLA Management

More is required than just knowing when a SLA has been (or is about to be) violated. Two separate aspects come into the more comprehensive type of SLA management: support for reacting to SLA violations and support for defining and managing SLAs.

Reacting to SLA Violations

The service provider should have an ability to automatically react to SLA violations[5]. Typically, when an SLA is violated, a number of actions need to be taken. These could be defined in terms of the business processes that need to take place.

The process that needs to take place could affect many different systems. Depending on the nature of the violation and the contract in place, it may be necessary to

- Reconfigure the service to restore proper operation
- Raise the priority of the network problems that are causing the violation
- Notify the customer of the violation
- Notify the customer's sales representative of the violation
- Notify the billing system for the customer to receive any credit due

Whatever the actions, these should be defined as a process in a process engine. When a violation does occur, the service-assurance system can then spawn an instance of the appropriate process.

Defining an SLA

The second part of managing an SLA is defining it. Though part of this process is a legal and contractual process that falls outside the scope of service assurance, there is a critical portion that needs to be addressed here.

Service-level objectives are what make it possible to measure SLAs. A critical portion of defining the SLA is defining the SLOs. These SLOs need to be well defined. It is important that the definition include the following three elements:

- The measurements—this involves determining what will be measured (e.g., the time it takes to complete one sales transaction)
- The conditions—this involves determining what constitutes an out-of-bounds measurement (e.g., transaction time over 500 ms)
- The method—this involves determining how the measurement will be made (e.g., a test transaction originating on a remote node will be measured every 10 minutes)

If these three variables are well defined for each SLO, then there should be no arguments between customer and service provider about meeting or not meeting the SLA. As ambiguity comes into these definitions, however, uncertainty and inconsistency between customer expectations and service-provider expectations stemming from the SLA will increase. Such uncertainty and inconsistency will often result in unsatisfied customers.

Clearly, once the SLOs have been defined, the service-assurance system needs to be set up to measure and report on them.

Business View

The business view allows the service provider to see how well their business is running. Rather than looking at individual services sold to particular customers, it looks at a service offering across all customers.

The business view links the service-assurance process with the other processes, along with the broader financial picture, to define cost accountability. The typical service provider has no good way to determine the cost of providing and operating individual services or types of services.

This paper presents the business view as part of service assurance. In reality, though, this view needs to cross all four of the business processes. It could equally well be described as fitting in service delivery, service usage, or CRM.

Service-Offering Management

Service-offering management provides the service operator with a business view across all instances of a particular service. Rather than just focusing on a service provided to a single customer, looking at a service from the business view allows the service provider to get answers to questions falling under the following topics:

- ***Return on investment (ROI) and profitability of service:*** Is the service offering getting us the ROI expected by our investors?
- ***QoS offering:*** Is the service offering meeting the quality standards that our customers expect from us?
- ***Ability to meet SLAs:*** Are we generally meeting our SLA commitments?
- ***Customer satisfaction and retention:*** What is the level of satisfaction of our customers for our service offering?
- ***Cost accountability:*** Are the proposed costs in line?

This is an area where only rudimentary work has been done.

Customer View

While the other three views provide information to the service provider, the customer view is information for the customer. There are three major functional aspects that need to be provided:

- Status update
- Trouble reports
- Service reconfiguration

In some ways customer views appear to be identical to service views. However, there are some important distinctions that benefit from it being considered as a different view.

As soon as access to parts of an OSS becomes available to outside users, security and access control to information need to become a primary architectural driver. For customer views into a service-assurance system, this translates into ensuring that the following security issues are addressed:

Protection between customers—customer A should not be able to see anything about customer B (their existence, the services they use, the amount of traffic they generate, etc.); this is even the case when the services provided to customer A and customer B share the same equipment (e.g., an edge router)

Abstraction of the network—customers should not be able to see the details of how a service is provided; they should only see the abstraction the service provider wants to present to them; this needs to be detailed enough for the customer to understand and control their service, yet abstract enough to allow the service provider flexibility to make changes and protection of their proprietary ways of providing a service.

Status Update

The first service-assurance aspect that customers want from a service provider's system is a view of the current status of their services. The customers need to be able to see their "logical network" where such a view is applicable, as well as to see overall service levels and service levels relative to their SLA.

Because of the security issues addressed above, it is generally desirable to keep the details of the customer view at a service or service-element level. Even so, it is generally not sufficient to simply provide the customer with a feed to the service monitoring or SLA monitoring information. Since these are geared as internal systems, they may not have the proper abstraction level or filtering of information to adequately protect the information of either the customer or the service provider.

Trouble Reports

The second aspect of the customer view is the need for customers to be able to submit trouble reports and track the progress of resolution. There are three considerations that become important.

First, though the customer needs to be able to track the progress of the trouble reports, there will be much information in the trouble report that may be inappropriate for the customer to see. The service provider could consider things like names of technicians, actions taken, or details of the service implementation as proprietary information.

Second, the customer should also be able to track the progress of select trouble reports that relate to the customer's service. If the service provider has already detected a problem, it should not be necessary for the customer to create a new trouble report just to be able to track the progress to resolution.

Finally, a well-architected system will allow for a mapping from the service on which the customer is reporting a problem to the components in the network providing the service. This will allow automatic testing of the components.

Service Reconfiguration

Finally, if we look at service assurance as the ongoing operational support of a customer's service, it also includes customer ability to reconfigure profiles and service parameters.

As systems implementing policy-based management (see the "Policy-Based Management" subsection within the "Future Directions" section for further discussion) become prevalent, the customers will want to change some of the policy parameters that apply to their service.

Process Integration

As pointed out before, the service-assurance processes should not be viewed in isolation of the other processes. This section provides examples of some of the interactions that should be expected (see *Figure 3*)

Meeting Service-Level Objectives

Since the service-assurance systems are responsible for ensuring that the contracted service levels are being delivered, all processes need to feed information to the service-assurance system about meeting or not meeting specified service-level objectives. As described in the "SLA Monitoring," SLOs can address such things as hold time while waiting for a CSR or time to provision new services. This information needs to be made available for a comprehensive view of SLA compliance to be possible.

Service Delivery

The service-delivery processes include all systems and steps related to the process of implementing a new customer account and activating a new service subscription. Actual service delivery is the result of a chain of activities that starts with a service request to a central point and ends with providing the requested services. Typical process steps in between include the following:

- Confirming the service request
- Identifying the type of requested service
- Checking if the required facilities are available
- Informing the billing system
- Activating the service
- Registering the subscription

Three specific areas of integration between these processes and service assurance deserve elaborated discussion:

- Network inventory
- Testing
- Service restoration

Network Inventory

Inventory and configuration management is an area that spans the service-delivery and service-assurance processes. In broad terms, the service-delivery process uses inventory management to understand what resources are available, plan the set of resources that will be needed to provision a service, and allocate the resources to the service.

Within the service-delivery processes, the network-inventory system has that knowledge and supports the following functionality:

- Equipment—the description in data format and graphical format of the physical layout of an NE (rack, sub-rack, cards, spares, etc.), allowing the service provider to plan and grow the equipment as needed to meet the requirements for new services
- Physical connectivity—the description of the physical connectivity between NEs, both in the central office (CO) and outside (connectivity to and from digital distribution frame)

FIGURE 3

Integrated OSS Solution

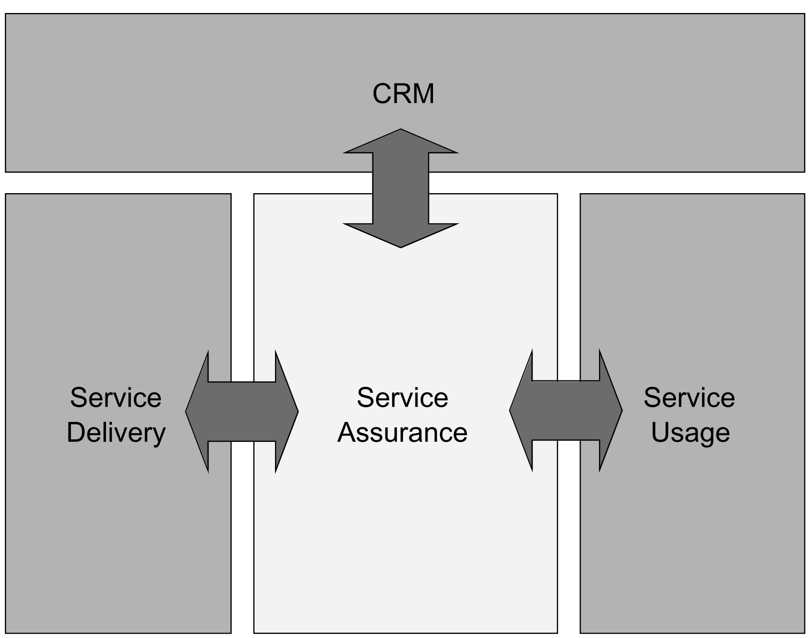

- Logical connectivity—the description of the logical connectivity; allocation of signaling channels, traffic channels, and operations, administration, and maintenance (OA&M) channels; and bandwidth allocation
- Services—the association of the physical and logical connectivity to the end-to-end services offered to customers as well as the association with the customer data

For service assurance, the inventory- and configuration-management function has been outlined earlier in the "Inventory and Configuration Management" section. This strong overlap should be addressed by a consistent system that addresses the needs of both service assurance and service delivery.

This consistency is more than just using one system for both processes. Creating the consistency is complicated because neither of the processes can be considered as the ultimate "master." Though the service-delivery processes are the ones that primarily create and modify much of the data in the inventory, the service-assurance processes can also modify it. In particular, as the service-assurance processes discover new components, they need to be able to reconcile the "as-is" network with the "as-designed" network.

Testing

Both service assurance and service delivery require a testing facility. Within service delivery, testing is used to ensure that the provisioned service works. It may also be used to determine if it is possible to provision a service (e.g., checking the line condition before provisioning a DSL circuit).

Service assurance uses testing to diagnose components and to verify components are operating properly.

Service Restoration

There are two scenarios for how service restoration can work. Both require the service-delivery processes to cooperate with the service-assurance processes.

First, the service could be configured with backup components. If this is the case, then service restoration principally consists of activating the backup components. For this to work, the service-delivery processes need to assign backup components and communicate them to the service-assurance processes.

However, in some situations it may not be desirable to assign backup components. In such situations, service restoration is still possible. However, instead of simply activating backup components, the service-assurance processes would need to request the service-delivery processes to re-provision the service using different components. For this to work, two things are necessary. First, there needs to be a way for the service-assurance processes to initiate a re-provisioning request. Second, there needs to be a way for the service-assurance processes to inform the service-delivery processes which components have failed and should not be used.

Service Usage

The service-usage processes deal with the collection of usage information from the network and its application to billing, call-behavior analysis, and other systems that can utilize this information. Systems that support these processes include the following:

- Billing mediation
- Billing
- Usage analysis

Two areas in which service usage interacts with service assurance are examined below.

Call Data Records Monitoring

Call data records (CDR) are records produced by the components in the network to detail the usage of the network. The initial purpose of CDRs was to allow the service provider to bill their customers on usage. However, since CDRs capture how the network is being used, they can provide additional valuable information.

In particular, by monitoring and analyzing the CDRs for predefined patterns, it is possible to generate alarms that indicate possible problems. For example, it is possible to identify a sudden drop in call-completion rates or an increase in congestion.

For this type of detection to be possible, the service-assurance systems need real-time access to the CDR information collected as part of the service-usage processes.

SLA Violation Credit

Often times, an SLA will contain penalty clauses that call for the service provider to credit the customer in cases of noncompliance. It is also conceivable that an SLA could be written that calls for a bonus to be paid if certain service-level objectives are exceeded. From a process and integration perspective, this involves a service-assurance process detecting the SLA violation (or exceeding the SLO) and requesting a credit be issued (or a bonus payment charged) through the service-usage processes.

CRM

The CRM processes involve those processes that provide the contact between the service provider and its customers. The goal of CRM is to allow service providers to better identify, attract, serve, and retain customers by managing customer interactions across multiple channels. Included within these processes are both traditional and new types of systems:

- Traditional call centers
- Sales-force automation
- Web call centers
- Web-based self service

Service and SLA Information for CSR

For CSRs to have the ability to successfully provide support to the customers, it is imperative that they have access to the right information. For a service provider, one of the critical pieces of information is the status of the customer's service.

The CSRs do not need to know the detailed status of the service components. The right level of information is the service status and the SLA status of the customer's service. Rather than expecting the CSR to utilize a different system to access this information, it should be supplied to them through their CRM system. This requires the service-assurance systems to communicate the information to the CRM systems.

Portal for Customer Views
To be successful, CRM needs to have complete understanding of the interactions between the service provider and the customer. The service-assurance customer view provides customers with an understanding of their services. It, therefore, stands to reason that the service-assurance systems need to provide the customer view through, or under the control of, the CRM systems.

Having a single portal from the customer into the service provider would provide this. One of the pieces of information available through the portal would be the service and SLA status.

Verify Service
Sometimes the customer might perceive a problem with their service when there is no indication that there is a service problem. This could either be because the problem is with the customer's equipment or because a service component failed without reporting a problem.

In either case, the customer or the CSR may wish to verify the service. Service verification provides a link from the CRM systems into the service-assurance systems. It allows the CRM systems to request that a service, or a set of service components, be tested to ensure that everything is functioning correctly. Though the request comes from the CRM systems, the actual tests need to be under control of the service-assurance systems.

Future Directions

Policy-Based Management
Policy-based management is an evolving technology that will allow components to adjust their functions based on pre-set policy. Initially intended for the IT environment, these policy-based systems, as they evolve, will begin to address important aspects of the service provider's infrastructure.

Policies, in one form or another, have been used for a long time to manage networks and systems. Policy-based management is developing new tools and systems to automate many of the tasks. A significant difference of these new policy-based–management systems is a centralized repository of the policies. This allows all components and management systems to share a consistent set of policies.

The policy-based systems are imbedded into the NEs, service components, or element managers. They allow the operator to define a set of policies to govern the behavior of the component.

Rather than eliminating the need for service-assurance systems, however, policy-based management has the potential for complicating service assurance. When everything is working properly, then the customer should be receiving the contracted-for service level. However, service-assurance systems are not as valuable when everything is working—they earn their keep when things go wrong.

Policy-based management will simplify some of the processes of SLA monitoring, but it will also introduce an additional failure point—the policies. As policy-based–management systems evolve, it will be necessary for service-assurance systems to evolve with them.

Federated Management
In today's world of complex communications services, a service provided to a customer is often composed of SEs purchased from another service provider. The TMF defines this as the layered-service concept, as illustrated in *Figure 4*.

In *Figure 4*, a customer obtains a service with one or more service access points (SAP) from Service Provider 1. Service Provider 1 integrates various SEs to create the service. One of these SEs is a service that Service Provider 1 obtains from Service Provider 2.

Two conflicting issues complicate these types of layered services. First, service management and diagnosis require an end-to-end view of the service. Service Provider 1 needs to have management-information flow across the boundary with Service Provider 2.

Second, business requirements restrict the sharing of information across administrative boundaries. Though in this particular case the two service providers are partnering in providing the service to the customer, they may also be competing in other areas.

Interconnection gateways that provide controlled access to ordering and trouble-ticketing systems are beginning to show up. These will need to be extended to provide a more complete solution. In particular, there need to be standard ways to do the following:

- Selectively share management information across administrative domain boundaries
- Derive measurable aspects from the legal SLA documents
- Provide recommendations and policies to define metrics and their bounds for service compliance

References

1. Elizabeth K Adams and Keith J Willetts, The Lean Communications Provider: Surviving the Shakeout through Service Management Excellence, McGraw-Hill, New York, 1996.
2. D. Awduche et al., "Requirements for Traffic Engineering over MPLS," Internet Engineering Task Force RFC2702, September 1999.
3. Preeti Bhoj, Sharad Singhal, and Sailesh Chutani, "SLA Management in Federated Environments," HPL-98-203, Hewlett-Packard Company, Palo Alto, CA, December 1998.
4. "Maintenance: Telecommunications Management Network, TMN Management Function," CCITT Recommendation M.3400, October 1992.
5. "Customer View of Internet Performance: Measurement Methodology and Metrics," Cross-Industry Working Team, Corporation for National Research Initiatives, October 1998.
6. Jeff Drake et al., "The HP IT Service Management Reference Model," document 5966-4808E, Hewlett-Packard Company, March 1998.
7. "Benefits of a Service Assurance Solution—Business Analysis," Hewlett-Packard Consulting White Paper, Hewlett-Packard Company, Grenoble, FR, July 1998.

FIGURE 4

Layered-Service Concept

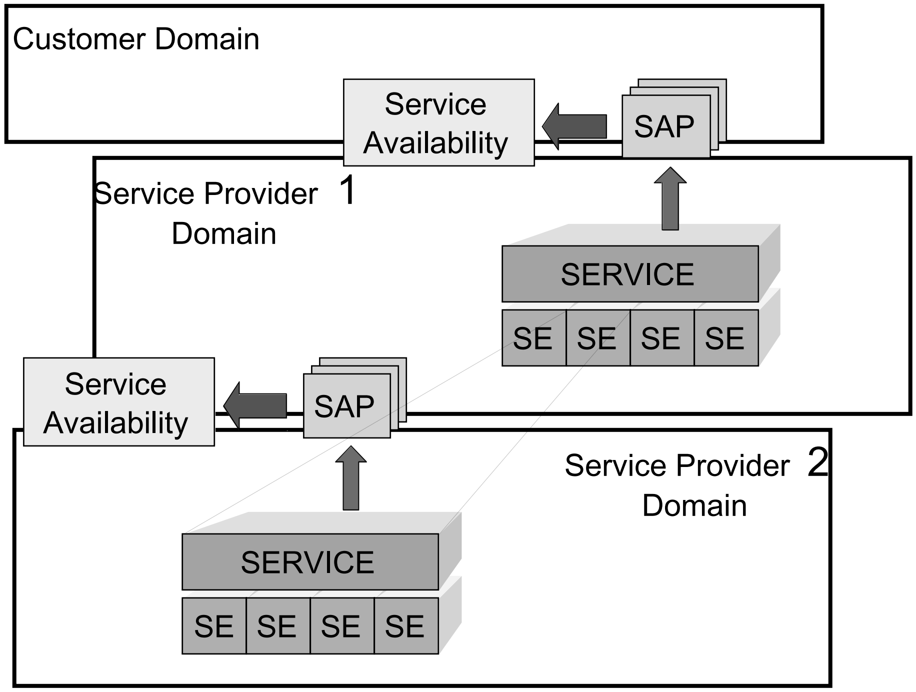

8. "SLAs—An Emerging Trend in the Internet Services Market," Internet Service Management Solutions Strategic White Paper, Hewlett-Packard Company, Palo Alto, CA, September 1998.
9. Dave Kosiur, "Policy-Based Network Management: What Is It and Who Needs It?" Network Strategies Overview, The Burton Group, Midvale UT, November 1999.
10. T. Li and Y Rekhter, "Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)," Internet Engineering Task Force RFC2430, October 1998.
11. "Performance Reporting Concepts & Definitions," TeleManagement Forum, TMF701, Evaluation Version 1.1, May 1999.
12. "SMART TMN™ Network Management Detailed Operations Map," GB908, Evaluation Version 1.0, TeleManagement Forum, March 1999.
13. "SMART TMN™ Telecom Operations Map," GB910, Version 2.1, TeleManagement Forum, July 2000.

Notes

1. In this paper, the term "service provider" is used to refer to communications service providers (CSP) or data service providers (DSP), including those providing value-added services, such as application hosting, unified messaging, etc. The acronym xSP can also be used to refer collectively to CSPs, ISPs, application service providers (ASP), etc.
2. It is called "network" view to match the well-known TMN model. This, however, is an unfortunate naming since it seems to exclude non-network components, such as platforms and applications. This is not the intent.
3. This is where the component fails without sending any notification of its failure.
4. Reported numbers range from 20 to 60 percent or more.
5. This, of course, could include automatically reacting at some threshold before an SLA is violated.

The OSS Integration Quandary

Ryan Licari

Independent Telecommunications and OSS Consultant

Introduction

Communications service providers (CSPs or carriers) are no longer measured by their billion-dollar investments in high-speed networks, but rather by operational excellence and customer profitability. Profitability, not the number of cities lit, or hitting the next subscriber target levels will keep carriers in business. How does a carrier then create an organization and operations that can sustain profits? The path to profitability hinges on a combination of differentiated service offerings and operational excellence—with automation and flow-through integration being the key to operational excellence. Fully integrating key components of a carrier's operations support system (OSS) is the most critical factor. The end goal is to drive manual interaction to next to zero. Carriers tend to underestimate the true size of a complete integration effort with dozens of systems, a number of business processes, and even more customer and transaction data. This paper addresses all types of carriers including incumbent local-exchange carriers (ILECs), competitive local-exchange carriers (CLECs), interexchange carriers (IXCs), metropolitan-area network (MAN) carriers, Internet service providers (ISPs), and broadband network operators such as digital subscriber line (DSL), fixed wireless, and cable operators. The paper provides and overview of the OSS environment today, an analysis of the true size and challenges of typical OSS integration, and an overview of the popular integration approaches utilized today.

Overview of Today's OSS Environment

To understand the true magnitude of the integration challenges facing carriers today, we must first gather an understanding of what a typical OSS looks like. It may be easier to look at greenfield carriers and more mature carriers, such as regional Bell operating companies (RBOCs), separately, since the integration elements will be substantially different. Analyzing the OSS elements of a greenfield carrier first, you will find the luxury of a small number or even no legacy systems in place. To start, these carriers will most likely evaluate and purchase a packaged customer-care and billing system from a vendor such as Amdocs, Convergys, Lucent Technologies, Portal Software, or CSG Systems. Billing tends to be the obvious first choice, because if you cannot bill a customer, a carrier will not be in business too long. Next comes order management and provisioning systems. These systems provide true measures of operational efficiency, enabling carriers to streamline their service fulfillment and service-assurance processes. Packages such as MetaSolv Software, Cramer Systems, and Sigma Systems for order management, and Syndesis, CoManage, and 3Com for provisioning are the common vendors mentioned in this category. Customer-relationship management, or CRM, tools soon follow. These systems enable easy management of a carrier's service offerings and the overall customer ordering and acquisition process. Siebel Systems is the lead horse in this race followed by PeopleSoft. The final primary tier purchase is usually fault management. It is essential for carriers to be able to easily identify network outages or problems. Failure to adequately do so can result in very costly customer churn from abnormal and lengthy network outages. Micromuse is the category killer here.

The next set of purchases usually involves what many carrier's consider as "nice to haves," selected in the second tier of OSS buildout. These areas include interconnection gateway, service-level agreement (SLA), and performance-management tools. In the interconnection gateway space, NightFire Software, Quintessent Technologies, Telecordia Technologies, and Accenture's Launch Now solution all compete. Interconnection gateway solutions enable carriers to streamline the provisioning process of circuits that reside off their own network. For example, a carrier may need to obtain a local loop, port, or frame-relay circuit from another carrier to fulfill a service request. Interconnection gateway software intelligently validates and automates these orders, eliminating errors and speeding customer turn-up. Service-level management solutions are often addressed through a variety of packaged and custom reports generated from other OSS packages. This approach may provide a certain level of monitoring and SLA management, but it is a very manual intensive and somewhat crude process. There are some companies, such as the European-based joint venture between Verizon and Telecom Italia called Sodalia Solutions, that are starting to successfully position their SLA management tool in larger accounts from wireless carriers to long-distance providers. By probing for data on the network and throughout other OSS systems and meshing this data with easily defined SLAs, this tool monitors, manages, and alarms SLA infractions in real time. Again, this can often be viewed as a "nice to have," but the benefits to the customer from proper implementation of these tools can be quite powerful.

Let's now move from the greenfield to the more complex scenario—the OSS components of larger carriers such as the

RBOCs that have been in business since the time when long-distance service was actually a money-making proposition. Telecordia Technologies and the big integration companies, such as the former Andersen Consulting and Price Waterhouse Coopers, have built from scratch enormous systems to support the expansive business requirements of RBOCs. Some of these systems are on record as being the largest systems ever created. Even the names of theses systems can be complex: CABS, CRIS, SPC SOS, and FlexCab. Because these systems cover incredible amounts of business functions, they tend to require constant code changes and updates to user interfaces. They are much more dynamic and expansive systems than the "off-the-shelf" solutions typically implemented by the greenfields. In addition, because the carriers have been in business for longer, they support a larger number of service offerings, each of which may have its own billing, customer-care, order-management, and inventory systems. The number of custom systems can grow exponentially.

The Magnitude of Integration

This section outlines a number of elements that make creating and maintaining successful integration a substantial effort. First, the number of systems alone can be daunting. The greenfield integration scenario appears to be more straightforward versus the integration tasks of more mature carriers' systems, but the number of systems requiring integration in either case can be staggering. The greenfield scenario must have a minimum of four and upwards of a dozen disparate OSSs to run their business, but the legacy carriers' can have hundreds. Let's consider a "killer" consumer offering: local, long distance, DSL, and a fax line all on one bill. Nice offering, but legacy carriers need to tie together the CRM, order-management, and provisioning systems to ensure an accurate service order and fulfillment. All of the services must then be tied together from multiple billing systems to be placed on a single bill. That is consumer only, not even mentioning the exponentially more complicated business offerings.

An additional complexity involved with integration is the actual method available to migrate data into and out of systems. The off-the-shelf packages have a leg up in this area. Off-the-shelf packages tend to design and build application programming interfaces (APIs) that allow for the export and import of critical information. These APIs were designed and built as part of the initial product development, which makes for better throughput and easier integration points. These APIs are intelligent and can massage information into the format required by the system, and they also inherently know how to update the systems to keep consistency and integrity within the application.

Unlike the packaged software solutions, the custom systems of the legacy carriers do not come equipped with readily built APIs to exchange critical data, requiring additional custom work to extract and input information directly into the systems core, usually with direct database writes. This method comes complete with risks. The systems are so complex that a single transaction posted from another system may require tens of table updates and database trigger modifications.

Looking at the type of data that needs to be exchanged can further help us to understand the complexity involved in systems integration. Some examples of information that needs to travel back and forth between systems, while maintaining integrity, include the following: customer information, billing information, service information, order information, and usage information. The most simple on the list is the customer information. This information is usually relatively standard across systems, so getting name, address, and contact information between systems is relatively simple. Billing information adds the complexity of payment type, terms, rates, and conditions involved with a customer's bill. Making a broad leap to the next level of complexity, one quickly encounters service information. Service information is stored in a product catalog. Sounds simple? Hardly. The billing, CRM, and the order-management systems all have their own product catalogs with information stored in a variety of different formats. One of the largest challenges in the OSS world today is to enable timely and accurate updates of a service change through a number of product catalogs. Order information presents a different level of complexity. The issues here are two: 1) What is the database of record for an order, and 2) orders can be extremely data intensive entities. These orders can be customer-to-carrier orders or carrier-to-carrier orders. For carrier-to-carrier orders, such as local service request (LSR) or access service request (ASR), the order information is traditionally entered into an order-management systems, but the information must be integrated and transferred through an interconnection gateway solution, where the order is transmitted to the receiving carrier. The database of record should be the order-management system, but what happens when an error occurs either with the gateway or the trading partner? The error information must be relayed back to the order-management system to maintain data and order integrity. This is another integration point that is very dynamic and data intensive. If a carrier executes 10 types of orders to eight different trading partners, that is 80 different interface points that must be maintained. Each interface, and ASR order for frame-relay service to BellSouth, for example, can contain as many as 10 message types, 20 forms, and hundreds of fields each. Keep in mind that all of these interfaces are not static, but rather change frequently.

Another key element to consider is the constant updates made to both off-the-shelf and customer applications. Packaged software packages tend to produce at least one major release upgrade per year, as well as a number of minor releases and bug defects on a consistent basis. Each change in a software package causes the potential for an integration change. Along with the potential code changes to integration components, carriers must also execute a new suite of integration tests to ensure that the systems are still compatible. As depicted, the versions and transaction data are never static and the wealth of data that needs to be migrated throughout systems can be daunting.

Looking at a real-life example can help to put the complexity in perspective (see *Figure 1*). A small business places an order for virtual private network (VPN) service between its headquarters and all of its remote sales offices from a local CSP. A customer-service representative (CSR) enters the order into a CRM package, where the account is created and pricing is negotiated. The service and billing information must be transmitted form the CRM to the billing and order-management systems, which breaks the order into its components. One of the sales office locations is in Washington, D.C., but the carrier taking the order currently lacks network capacity in that region. The network capacity was determined by transmitting order information from the order manager to the inventory system for validation and back. Satisfying this order now requires the carrier to order circuit capacity in Washington. The off-network circuit request comes in the form of an ASR entered in the order manager and sent through and interconnection gateway and on to the Washington, D.C.–area trading partner for provisioning. The responses from the trading partner must flow back through the gateway and on to the order manager. Now the order manager sends service and provisioning information to the provisioning and activation systems to turn up the service. The complete VPN service can now be turned on, leaving the fault-management system to transfer usage data and faults back to the order manager.

Potential Integration Approaches

We have now seen the business requirements and goals of a carrier driving the need for a number of integrated support systems, and we have discovered that integration itself is no lay-up—its time to explore the potential solutions of how integration occurs:

1. Multiple Data Entry
This option should not even be discussed. It's basically ignoring technical integration and using manual double entry. Carriers implementing this approach are asking for large-scale errors and service delays. The risks, challenges, and excessive expenditures surrounding this approach are obvious.

2. Custom Integrations
Custom integrations involve contracting with a systems integration firm to build "messaging tools" between systems to move data around the OSS. Custom integrations have typically been long, arduous processes, with no guarantees of success. This approach is shrinking in popularity.

3. Enterprise Application Integration (EAI)
Often referred to as middleware, EAI solutions offer tremendous functionality to automate business processes between systems, with very powerful tools to create and map critical business processes. This approach has been quite popular among carriers, but actual field implementations have yielded less than adequate results. EAI solutions can add a great deal of additional complexity to the integration, including adding another dynamic system that must be integrated and the need to maintain experts on whatever EAI tool you may choose.

Giving carriers, integrators, and vendors a great deal of credit by learning from their experiences of the boom time in telecom, a concerted effort is being made to create a hybrid integration solution that holds promise. It's a hybrid of the custom integration option, the EAI option, and a third option that will be discussed next: the point-to-point adapter. This solution provides carriers with the tools that EAI packages provide, along with a set of pre-packaged components and methods that will enable communications between selected packaged OSS software solutions. For example, webMethods offers pre-packaged components that enable Siebel System's APIs with Portal Software's APIs, but the specific data elements that require transport are not included. This leaves the need for continuous customization of the components using the EAI tools to meet

FIGURE 1

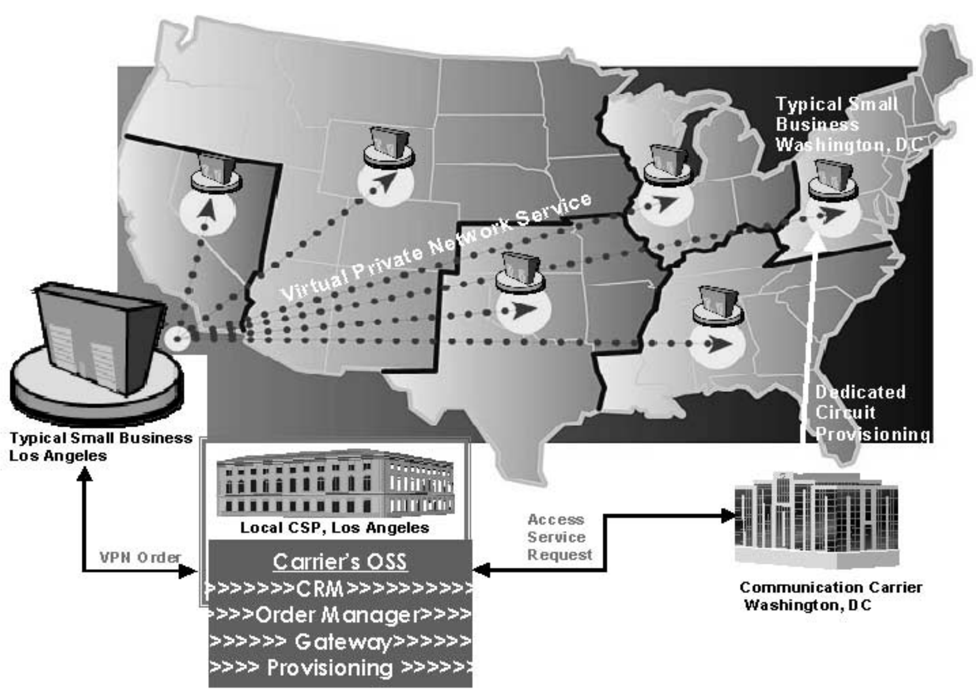

the needs of the carriers' specific business processes. EAI vendors are also going a step further in some cases by offering not only package-to-package components, but also a set of business processes or transactions to meet specific business requirements. For example, Vitria Technologies provides a set of customizable process flows and transactions to enable broadband service fulfillment processes. The systems integrators have taken a liking to this type of offering, creating large practices in a number of EAI technologies. The integrators have realized that integration using the EAI tools is more manageable and reusable that a straight custom approach.

4. Point-to-Point Adapters

Off-the-shelf, point-to-point adapters that connect and manage the communication of transactions between systems are starting to be offered by independent software vendors (ISVs) themselves. Point-to-point adapters offer carriers an additional, more headache-free option that helps to avoid the custom-integration pitfalls. These off-the-shelf, proven, integrated solutions can truly help carriers to achieve the data flow-through for which they are searching. Point-to-point adapters completely package integration. They contain specific messaging and process elements to manage specific processes between applications. The adapters also release a number of different versions to maintain productized support for updated versions of the packages themselves, as well as any changes in the data content or process flows between the systems. This sounds like a brilliant solution, but it is not one without downfalls. Chances are that ISVs will not offer productized connectors between all of the applications that you need to integrate. This approach also requires the ISV maintaining the adapter to become an expert and to establish incredibly close ties to the ISV to which it has built an adapter. This is an option to keep an eye on in the quarters to come.

5. One-Stop Shopping

The final option for carriers makes the finance department ecstatic because it's easy on the budget. However, the operations people may frown on the less than complete functionality. Companies such as BlueSpring Software, Wisor, and Telution are building all-inclusive solutions that address a carrier's billing, customer-care, CRM, order-management, and provisioning needs all in one. This option eliminates the need for integration and usually can be procured for a fraction of the price of buying the best-of-breed solutions in each area. The pitfall here is that the functionality of these packages may be more limited than the category killers.

Conclusion

Today's telecommunications climate is like a summer in San Francisco—changes are frequent and sometimes drastic. OSS components need to keep up with these changes. As we have seen throughout this paper, integrating solutions can provide carriers with powerful flow-through automation, enabling the ultimate in operational excellence. However, the path to utopia is a minefield. Although there are many different approaches that carriers can choose to integrate their systems, a few things are certain. Whatever integration method a carrier may select, an appropriate solution must be "future-proof"—i.e., technologically advanced enough to quickly adapt to change. Carriers should also be aware of, understand, and prepare for the magnitude of proper integration.

Defining "Carrier Grade"

Clyde Miller
Vice President
Nortel Networks

Introduction

The telecommunications industry has always prided itself on its reliability and availability, but the telecom network is currently undergoing a rapid state of change, as more services are being added to existing networks, new networks are being added to include those services, and new elements are coming into the network from nontraditional sources. And while the integrity of the network is still the number-one priority of the industry, competitive, first-to-market pressures are rising.

One of the results of these conditions has been the emergence of a wide variety of interpretations of the term "carrier grade." Without a firm consensus within the industry about what exactly constitutes carrier-grade products, services, and networks, the definition of this term will remain unclear and elusive. This paper discusses some of these interpretations, addressing the implications of leaving the term without a precise, standardized definition.

Carrier Grade—Why Care?

These new network elements (NE) work in performing critical functions and include new switching fabrics, new edge interfaces, and other powerful devices with sophisticated hardware and software. Many of the elements being introduced into networks are more related to a computer-type environment than what is normally considered a telecom-type environment.

Maintenance windows, as far as the network is concerned, are essentially disappearing, making the reliability of those elements even more important. The outage impacts of some of these new elements are quite significant, especially when a company is trying to introduce a leading-edge service to consumers and promoting regular use. In such an instance, up-time is paramount.

The industry lacks clear, specific, and consistent carrier-grade requirements, and products at the forefront of this transition often lack standardized guidelines. In light of these deficiencies, the first step is to define carrier grade; the next step is to explore how true carrier-grade service can be attained; and the third task is to determine the industry's level of responsibility for achieving it.

What Does "Carrier Grade" Mean?

As illustrated in *Figure 1*, carrier grade is founded upon four basic principles or "pillars": reliability, upgradeability, maintainability, and supportability. Without all four of these pillars, the overall integrity of the network is compromised. These capabilities, within any element of the network, have to be maintained.

Carrier-Grade Reliability

Carrier-grade NEs must be designed and manufactured to ensure robustness and longevity. Products and services must demonstrate the designed levels of performance measured against a uniform standard, such as TL9000. TL9000 is the evolution of the requirements and quality measurement system (RQMS), a measurement standard that has been used in the industry for some time. TL9000–based metrics provide data on outage downtime and frequency.

In terms of consistent availability or consistent objectives, the goal is for 99.999 (five nines)—a number that equates to about two minutes of downtime per year, a mean time between failures of 17 years, and a mean time to repair of 16 minutes. Both the network and the services need to be measured according to such standards. The first step in doing so is to break the network down into its component parts and measure the individual parts. The next step is to roll all the components up into the network and measure at that level. Finally, in some cases, a service-basis examination is necessary.

Achieving designed-in reliability requires operators to pay close attention from the very beginning to hardware- and software-component selection, manufacturing processes, environmental robustness (how the component operates in field conditions), and industry standards.

Carrier-Grade Upgradeability

Carrier-grade NEs must have the ability to flex with changing business needs and technology without impacting the end user. Software elements must allow seamless on-line growth, meaning they can accept feature additions as well as patch, generic, and upgrade applications. Such flexibility allows companies to develop hardware without out-of-service downtime.

FIGURE 1

What Does "Carrier Grade" Mean?

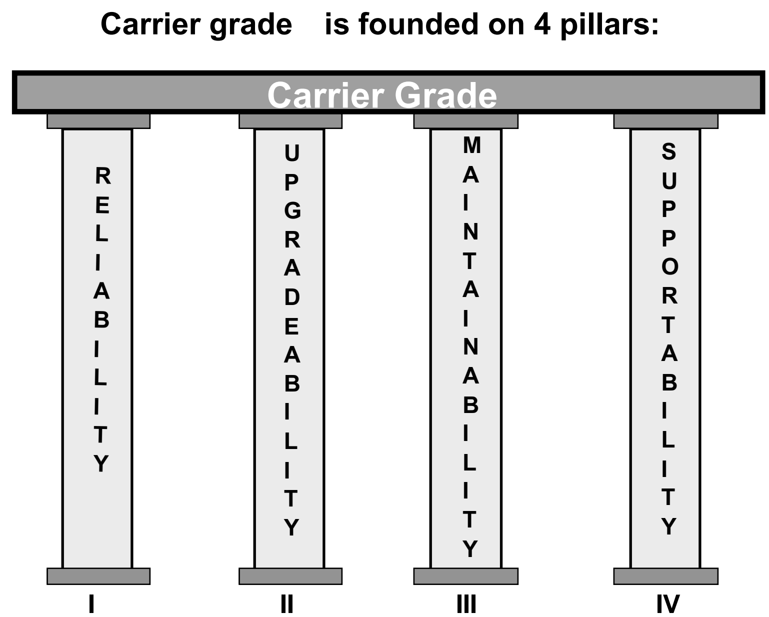

Hardware elements must allow for expansions, reassignments, and upgrades to the equipment. The hardware must also allow for network reconfiguration. For example, if a company needs to reroute traffic for service reasons, the system should allow movement to another element or another unit. The hardware must also have a uniform change process with adequate documentation.

Documentation is often written by scientists for other scientists. The problem is that most of the time the people in the field are not scientists. For that reason, the documentation needs to take more of a human-factors approach and to better address the ways the non-scientist thinks and does things. In the end, good documentation will help reduce procedural network outages.

Carrier-Grade Maintainability

Carrier-grade NEs must also have the ability to communicate easily with maintenance centers and to automatically perform a wide variety of maintenance activities. Self-diagnostics are important, giving the network the capability of routinely looking at itself and reporting if the operating system is functioning properly. Auto-imaging is another capability that needs to be built into all these elements because it systematically takes a view of the network. Auto-imaging can provide information on exactly what the network looked like before a reconfiguration of services.

Maintainability also includes on-line data changes and subscriber changes, system audits, auto-recovery, and auto-correction. The idea is that the system should be able to recover automatically when it detects an error through its self-diagnostic procedure. With auto-correction, the system should be able to reroute automatically if necessary or to change activity of its duplicated switch process or power supply.

Centralized access is also necessary for maintenance provisioning, monitoring, and usage. As the industry moves forward with this new mix of elements, people must be able to look at a total view of the network, understanding clearly what the elements provide as far as the end user is concerned.

Carrier-Grade Supportability

Carrier-grade vendors must have the mechanisms in place to support the installed base and meet industry standards for responsiveness. It is important to measure all the hardware elements and the responsiveness of the people on-line. One of the variables that TL9000 measures is how long it takes the carrier, service provider, or equipment provider to fix problems in the network.

Critical outages must be resolved within 24 hours, 100 percent of the time. Ninety percent of major problem reports that may affect service are handled within 30 days. Minor issues fall within the six-month window. These standardized metrics for TL9000 not only provide a framework for measuring outages but also provide a way to present the data that eliminates the ability to change scales or other variables to make a company look better than it actually is.

Furthermore, the people on-line must be technically skilled, available on a 24/7 basis, and available on-site. For a service organization, those issues are vitally important, as customers are rarely pleased when they hear that the technicians will look at their problem when they come to work the next morning or afternoon.

Carrier-Grade Network Element

For an NE, alarms are reported failures. For a typical carrier-grade NE for switching, the two minutes of downtime per year is equal to one incident per year. Availability is 99.9999. The element meets network-equipment building standards (NEBS) requirements and meets physical and human-factors design requirements.

This human element is important because although hardware and software are increasingly becoming more reliable, procedural errors are still occurring. People making mistakes is a growing concern in the industry, and such errors present the challenge of finding out what actually happened. Operators need to focus their attention on the way products are designed and the way they should present themselves via a human interface. The idea is to mirror how people look at processes. Product tests often show that people do not necessarily understand the element when they are asked to find a certain part in the frame or to find a certain element.

Additional requirements for a carrier-grade NE include the ability to self-identify its failures greater than 95 percent of the time so that people can take action prior to the product actually failing in service. Technical support must be provided on a 24/7 basis by skilled technicians. Redundancy is built into the design. Finally, the element must conform to zero outages even when reconfiguring the network, applying patches and maintenance releases and changing generics.

The zero-outage concept is relatively new. The industry formerly had planned outages. Nowadays, though, the system, ideally, should report a problem, and the on-line people should correct the problem, then reconfiguring to provide the service back to the individual as needed.

How Is Carrier Grade Obtained?

Failure within any one of the four pillars—reliability, upgradeability, maintainability, or supportability—should take away the ability to claim carrier-grade service. Elements that are not carrier grade do not permit reconfiguring, rerouting, or reloading without an outage. Once the network is in place, it is too late to consider these factors.

Carrier grade must be designed into the product itself. It cannot be added after the fact. *Figure 2* shows the overall product life cycle.

In the manufacturing phase, important considerations include the types of components and connectors selected and the types of processes used to actually build the components.

For the installation phase, grounding and earthquake resistance can be important. The key question is whether the element is resistant and stands up to the overall charges for that product.

During the upgrade phase, the product must permit processor changeout, hardware and software upgrades, and network-element visibility.

Maintenance considerations include the capability for the element to patch, auto-image, and auto-patch itself. Maintenance releases need to flow through the network. In terms of human factors, the more processes that can occur automatically, the less likely something will go wrong. Alarms and logs are necessary for maintaining the unit.

In the reconfiguration phase, the product must have zero downtime impact and software tools to help people walk through the configuration without any errors. If a system is designed properly, it is somewhat like a salt shaker. A salt shaker is a very simple mechanism: if used properly, salt comes out; if used improperly, nothing happens. The interrelationship of the reconfiguration is also important—knowing the effect that one action will have on all the other elements in that network.

Finally, in the retirement phase, hazardous material is a serious consideration. In the future, when the part is taken

FIGURE 2

How Is Carrier Grade Attained?

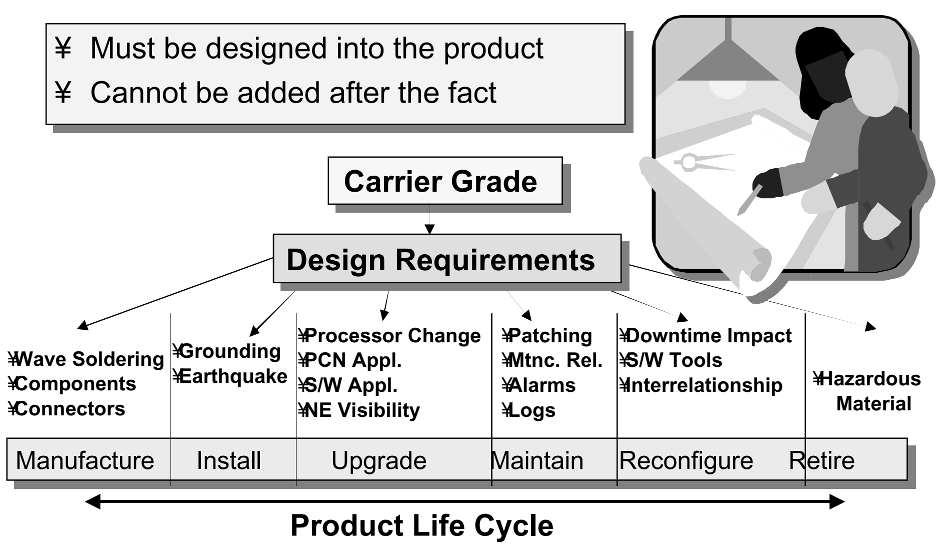

out of the network, disposal could potentially be compromised by the material that comprises the part.

What Is the Role of Industry?

The challenge to the industry is to establish formal carrier-grade requirements. Some standards are already well accepted, such as TL9000 and the evolution of RQMS, which is a global type of standard. Consequently, TL9000 will work anywhere in the world. Correlating traditional types of NEs to that standard is easier than working with some of the newer elements; so work remains to be done in the standards area.

These requirements need to be built on an element-family–type basis and viewed in terms of the service impact of the elements, the redundancy of the network, and other applicable pieces. As fabrics change, NEs on an individual basis probably become less important, but they are, and will remain, important as far as the network is concerned.

Once established, industry requirements need to be formally published so they become part of the decisions that drive the industry. Providers have to track the data to make sure that the elements actually meet the standards and to help drive continuous improvement in certain areas.

A few years ago, the standard for measuring outages was two minutes. Now, the bar is at 30 seconds and heading toward zero. It is vitally important to keep evolving the network and the elements themselves.

Adherence to requirements should drive business decisions on a day-to-day basis, always with an eye toward moving the standard forward or raising the bar as far as performance is concerned.

All these capabilities cost money, but a serious consideration of carrier grade involves understanding the real cost in the long term. After installation, a lower-cost product could end up costing much more than a higher-priced, carrier-grade product in service cost and customer complaints.

As shown in *Figure 3*, if products conform to the four pillars, they deserve the carrier-grade designation. Buyers of such products will be confident of being able to provide carrier-grade service to customers.

Conclusion

Carrier-grade service is increasing in importance. With data networks and services, there are now zero barriers preventing customers from changing from one service provider to another: customers can simply log on, cancel with one provider, and obtain service from another one. Service level, then, is going to become a paramount issue for providers, the linchpin for retaining customers.

FIGURE 3

The Four Pillars of Carrier Grade

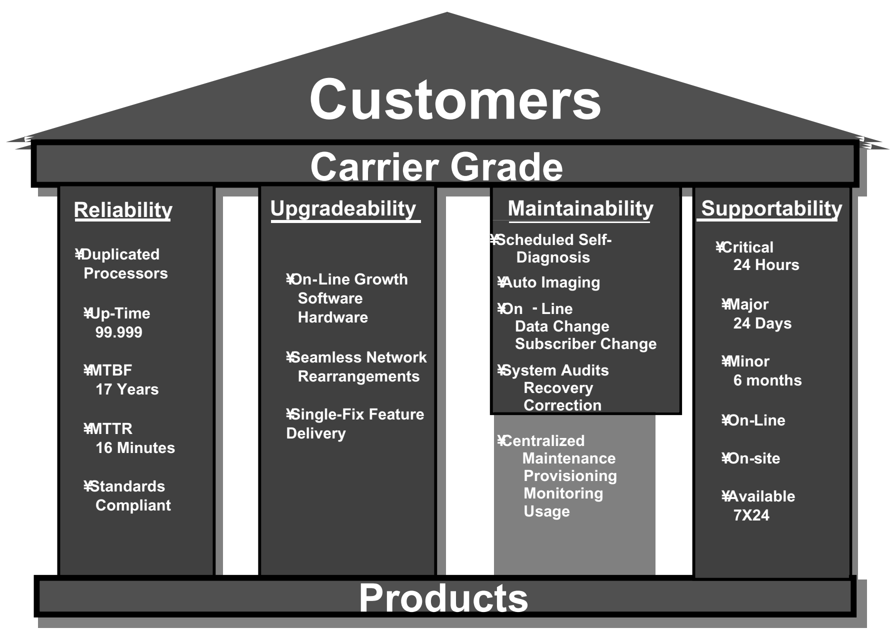

Process Flow in Service-Delivery Solutions

Richard Nietubicz

Practice Manager

Hewlett-Packard Consulting

Introduction

Companies are being challenged today by market conditions, business conditions, and enhancements in technology. To develop an operations support system (OSS) environment that is flexible, scaleable, and manageable, companies are reexamining their OSS strategies. There are several different types of OSS architectures: they can be common object request broker architectures (CORBA) or message-based architectures. This paper focuses on real-world implementation of OSS environments utilizing a message-based approach.

271 Testing

One of the most important test procedures in this area is based on Section 271 of the Telecommunications Act of 1996. Section 271 mandates rigorous and detailed validation tests on seven different processes: order management, inventory, trouble ticketing, network management, mediation, customer-relationship management (CRM), and billing.

The testing approach consists of three key areas: process definition, architecture, and implementation. Process definition, for which many tools are available, includes use cases and processes, activities, and sequence diagrams. Architecture encompasses integration events, rule-based logic, application logic, and object models. Implementation provides more information on the actual topography: client server, load balancing, monitoring, and system management.

The four regional Bell operating companies (RBOC) have shown particular interest in Section 271 because of their desire to enter the long-distance market. To initiate the process, the RBOC petitions the state commission and asks to have an independent test run to see how the company meets the 14-point checklist of Section 271. Reading the Telecommunications Act provides an entirely different perspective on what companies are supposed to do. In particular, the 14-point checklist mandates, among other requirements, an open and competitive environment; open and equal access to the network; back-office systems that are open, fair, and useable; and competitive pricing.

271 Testing Challenges

Section 271 testing involves an auditor role. Based on a statistical model, a test manager specifies the tests. Testing also involves a technology piece. Technology partners in such a test can create a pseudo–competitive local-exchange carrier (CLEC) with a fictitious name, a fictitious tax identification number, an operations center, and a marketing arm. The set of applications must meet the test criteria but may not need to be fully functional. Such arrangements can be made through the Federal Trade Commission (FTC), which supplies information and a special charter.

The fictitious company operates just as a real CLEC does, but with one key difference. A CLEC start-up must establish interconnection agreements and a gateway as quickly as possible to begin signing customers. In the pseudo–CLEC business, the testing company strictly follows the RBOC's published documentation. For example, if the steps are sequential, the pseudo–CLEC follows the steps exactly according to the sequence. If the step says that the RBOC will notify the CLEC, the test manager waits for the contact.

The pseudo–CLEC processes must include steps for sending information to and receiving information from the test manager. Another requirement is detailed logging of all interactions with the RBOC. Systems must be integrated across multiple companies.

271 Testing Implementation

The first implementation step is to start with the basic CLEC processes, such as order management and inventory, and to define additional steps for the exchange of information. Interconnection agreements have to deal not only with what products the pseudo–CLEC is going to sell but also with every product offered by the RBOC. Testing in this way provides detailed information about available services, profitability, and glitches in the process.

Twenty or 30 years ago, the standard timeframe for getting a service operational was between two and four weeks. Today, the timeframe is still the same. Because of complications in the back-end systems and across company process, new products can take even longer to rollout.

The real crux of the test is the logging system. The pseudo–CLEC must track how long each transaction takes. Part of this process involves spelling out duties, roles, and responsibilities explicitly. For example, the test process must define exactly what the call center is trying to do in order to identify problems with the technology. In many cases, the technology may not match the processes. Workarounds are common in the industry, but the questions to answer are these: How much time do the call-center workers spend in workarounds to execute a process? And how many errors are generated during this process? To rectify errors, the first step in the solution is to determine if personnel know about the process, if they are following the process, and if they are following it through to completion. If personnel are skipping steps, gaps in the process will inevitably result.

It is also necessary to establish procedures for a multicompany environment that involve real transaction screens and information exchange.

Telecom Processes

The integration of business processes has several different aspects: product, function, organization, and technology.

In terms of product, an integrated service manager (ISM) means many things to different people. For example, companies trying to write a business plan for a digital subscriber line (DSL)–ISM–integrated process generally do not look at it from a product base. This is a mistake, but it is driven by the nature of the business and by timing issues. Many architectures are available, but reality determines what the company must implement in terms of up-front costs and profitability.

In terms of function, a common practice is to establish an integrated call center to take care of Internet users and fold them into the same well-defined call processes.

In terms of organization, problems could arise from geography, power struggles, differences in tool sets, or any number of other sources.

In terms of technology, the principal issue is integrating broadband, voice, and Internet despite multiple mediation platforms and switches.

Order-Management Challenges

Process integration for order management presents several challenges. The first is order-type differences in which activities, such as customer-service requests and network-service requests, may be executed in different ways. By implementing a method for tracing through these specific scenarios, companies can determine if additional steps are necessary.

Validating the test when putting the process together is the key. After implementing the product-specific aspects, the timing of the information might be different depending on how the data values work in the scenario. Even though a process looks the same, it might not be the same.

Process and Application Integration

Figure 1 shows the scenario tracing for Levels 1 and 2 of order management. It is important to look at how the applications apply to a process and determine whether or not to go business layer or application.

The issue is whether to "adopt or adapt." What is missing is the product-level information or the data aspects. The test-

FIGURE 1

Order Management

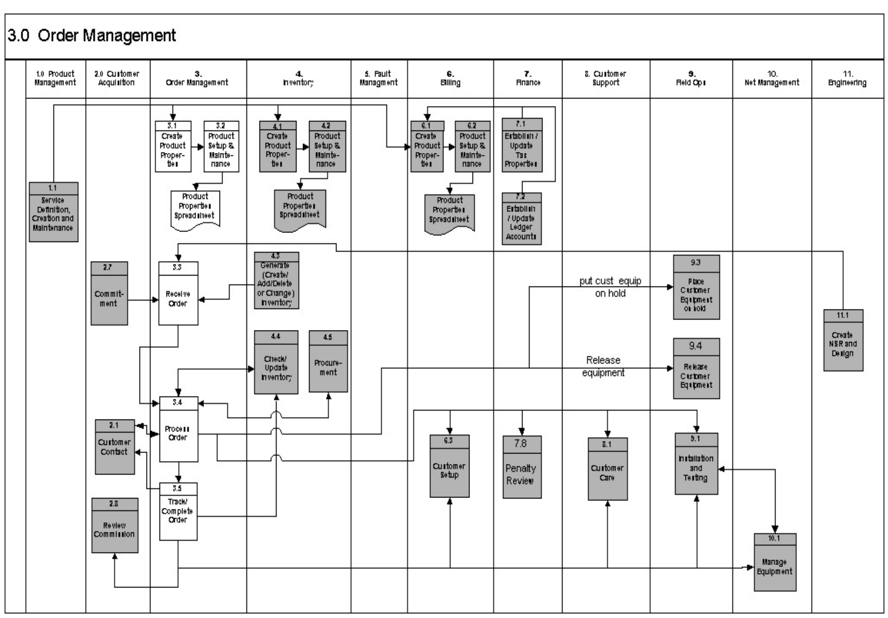

ing must include validation that will highlight differences between, for example, a customer-service request and a network-service request.

Figure 2 shows the scenario tracing for a network-service request, and *Figure 3* shows a similar diagram for adding an account. At this level of detail, companies can identify a number of problems: for instance, they can determine that the application does not meet the process, the data does not match the standard model, or the application programming interfaces (API) do not fit in the right place. Each application has its own implemented business process, so gap analysis must be performed between the business process and the application process.

Each application has its own data requirements and definitions. For data definitions, a popular solution is a common database. Implementation, however, may be precluded by costs. An alternative is data mapping between applications. Companies can develop an architecture or an implementa-

FIGURE 2

Receive Order (Network-Service Request)

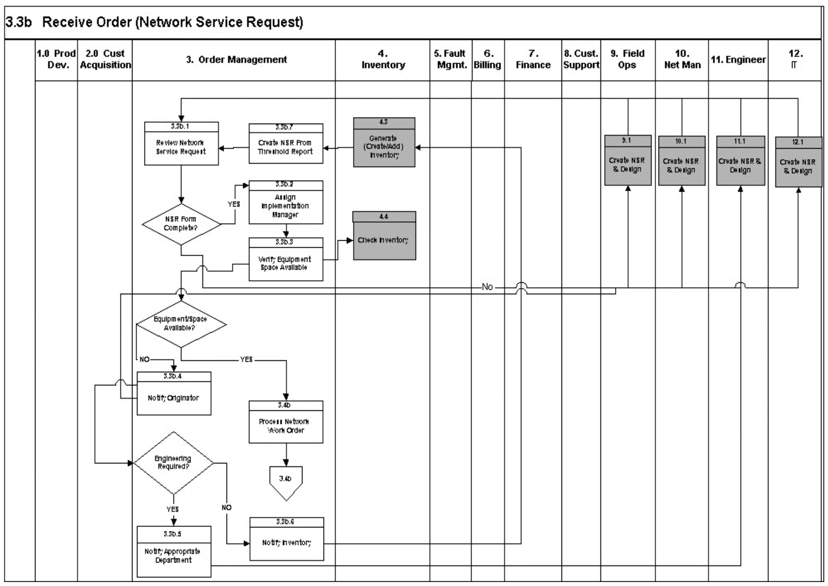

FIGURE 3

Add Account

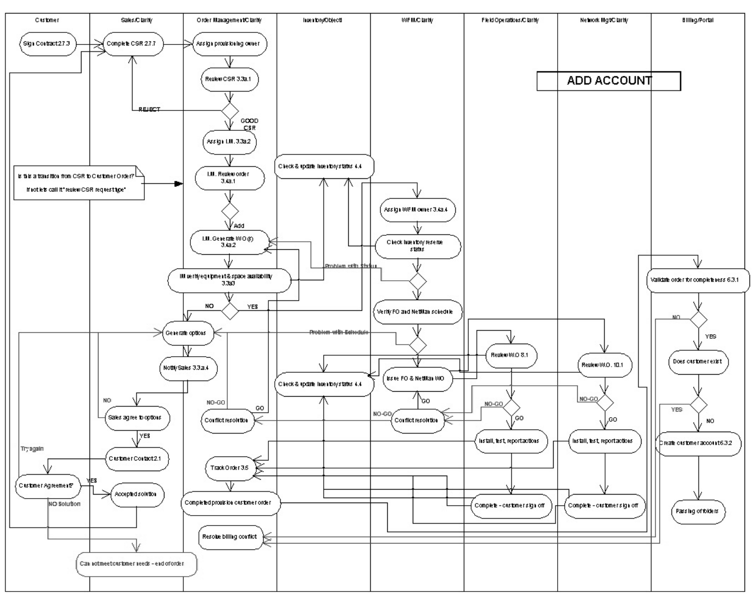

FIGURE 4

Integrated Service-Management Application Architecture

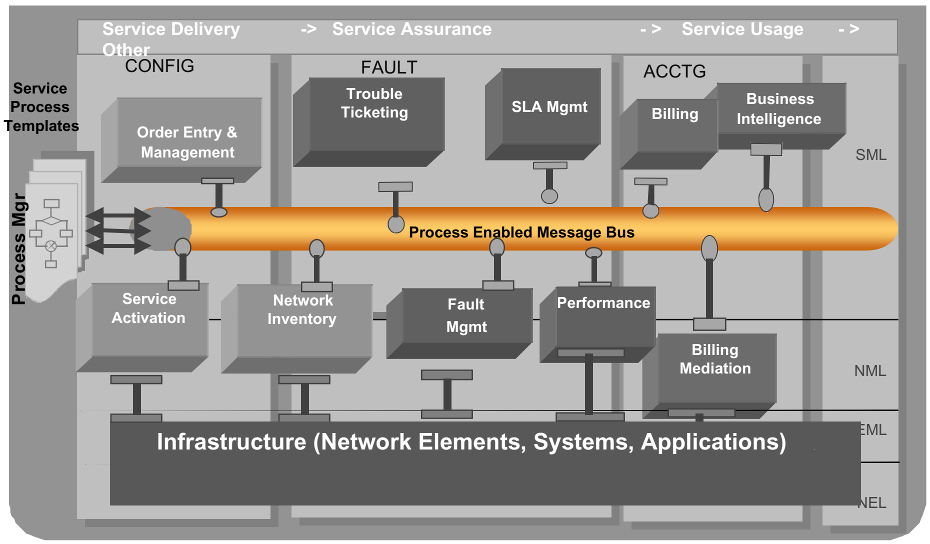

tion strategy that allows isolation to the data, first, and the capability of adding on, later. The product set is more important than the data model.

Integrated Service-Management Application Architecture

Standard architecture is common, but the issue is how to migrate to it. The solution is to start out slowly and choose a structure that will allow for growth (see *Figure 4*). Scalability is the most important factor to consider here, since new applications and processes will be a given. With the right tools in place, then, the company has a defensible baseline.

Environment Integration

Each application is normally implemented on its own configuration. The challenge is to architect an environment that supports an integrated messaging view of all applications. Each application also normally has its own file structure and database. The addition of a messaging layer means that an additional database must be synchronized and configured.

FIGURE 5

Environment Integration

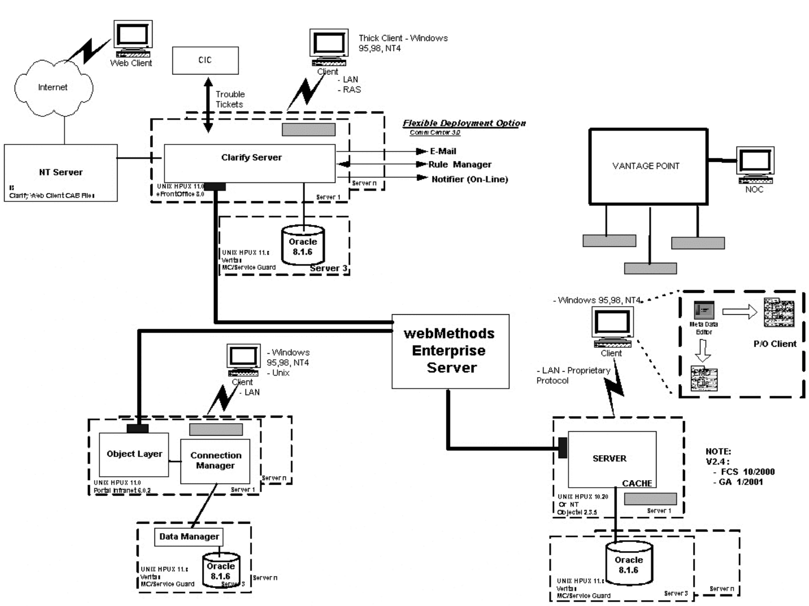

With each application residing in its own environment, the challenge is to ensure that the messages will flow between different system environments, including logging, backups, and file structures. Messaging environments tend to have their own monitoring capabilities. The challenge is to include these tools on the network-management desktops.

From an environment standpoint, legacy systems ensure that companies will never shed their old baggage. There will always be some signaling system 7 (SS7) need, and there is also the issue of next-generation versions one, two, and so on. Once a customer is on a specific system, the company is tied to it for a minimum of five years. One way to structure the architecture, as shown in *Figure 5*, is to modularize it and conceptualize at it as a set of parts, instead of looking at it holistically as a large path called order management. The modular approach allows additions and changeouts to the architecture. In a migration strategy, the company can define a new bundle or packet, add it, and then test it. If it works in that specific area, it could be applied to another area, allowing the company to drop in a new application and phase out an older one.

What Are the Gains?

On the business side, this effort provides a clear definition of business processes and procedures that establishes a baseline and results in user acceptance and buy-in.

From the system standpoint, companies gain an understanding of the capabilities of purchased software, thereby helping to minimize custom code and gain the ability to combine bundled items in a flexible environment. The result is a migration strategy and a purchasing strategy with a measurable return on investment (ROI).

From an implementation and investment standpoint, companies gain the ability to manage change through a phase-in approach and to optimize resources.

Lessons Learned

Companies can benefit from five significant lessons resulting from Section 271 testing:

1. Customer involvement is important. Users are needed throughout the definition of the process.
2. A phase-in approach allows the project team to refine the approach and ensure customer acceptance.
3. The application and messaging teams need to be involved from the beginning.
4. A clearly defined approach ensures consistency, completeness, and understanding of the processes and system implementation.
5. The change-management approach must allow for each change to be evaluated and reviewed before implementation.

Complexity of New Network Architectures Creates SS7 Management Challenges

Kim Parker
Product Manager
Tekelec

Introduction

Network operators often lack sufficient information to effectively manage their signaling system 7 (SS7) networks. Their primary source of information comes from the network switch(es) and a variety of basic instruments. Unfortunately, the usual source of the missing information is subscribers complaining about network problems.

With the addition of emerging technologies, such as Internet protocol (IP) transport, the integration of multiple protocols in a single network, the rapidly changing and currently unstable nature of IP–based signaling, and the constant introduction of new telephony services for both wireline and wireless networksæwith the addition of all these, SS7 networks are becoming increasingly complex. In turn, these changes make managing SS7 networks much more complex. To minimize the impact of new architectures and services and to manage the growth of the network, operators need advanced network-monitoring and -maintenance systems that provide real-time information, a platform for troubleshooting, and an information "feed" for business applications.

Challenges of IP Signaling

In a converged network, operators are often "blind" to the transport of messages over IP. No equipment exists to manage high-level telecommunications signaling over IP networks. In a converged network, operators often are unable to detect faults in the services provided by IP–based transport protocols, such as the following:

- Transport adapter layer interface (TALI) and media device transport protocol (MDTP)
- IP device control protocols, such as media gateway control protocol (MGCP), session initiation protocol (SIP), and session description protocol (SDP)
- Interworking between SS7–based and IP–based signaling
- Gathering and consolidating traffic statistical data across both SS7– and IP–based networks

As a result, with no monitoring system for the IP links, or for times when two systems are employed (see *Figure 1*), an operator does not have a consolidated view of the calls crossing the network. For example, when troubleshooting a problem, a customer is not able to perform a call trace that captures the entire call. Instead, an SS7 monitoring system shows the portions of the call transported via SS7, and another system captures the messages transported over IP. Additionally, telecom service providers are challenged by the many different protocols at work in their networks. Invalid interworking between protocols and interworking failures between service providers are at the root of network problems.

Today, operators must troubleshoot network problems utilizing various protocol-specific tools. For example, when tracing a call that deals with interworking protocols, operators must employ a separate view specific to each protocol. In general, there is no way to consolidate the messages traversing the entire network. A mechanism is needed to allow operators to visually trace a call that traverses multiple protocols within a single view.

Finally, each new service results in modified protocols, the enhancement of software loads, and the deployment of new equipment into the network. Each introduction of software and hardware into the network opens the door for new problems. By being proactive and identifying problems before they occur, or upstream, operators can incorporate these new services by verifying and validating the pieces prior to full deployment. Therefore, operators do not have to rely on customers to be the "eyes" of the operator.

The Solution

The best way to ease the burdens faced by telcos is to provide a single monitoring system for the entirety of converged networks (see *Figure 2*). Such a solution enables wireline and wireless service providers to monitor, main-

FIGURE 1

IP–Based Carrier Network

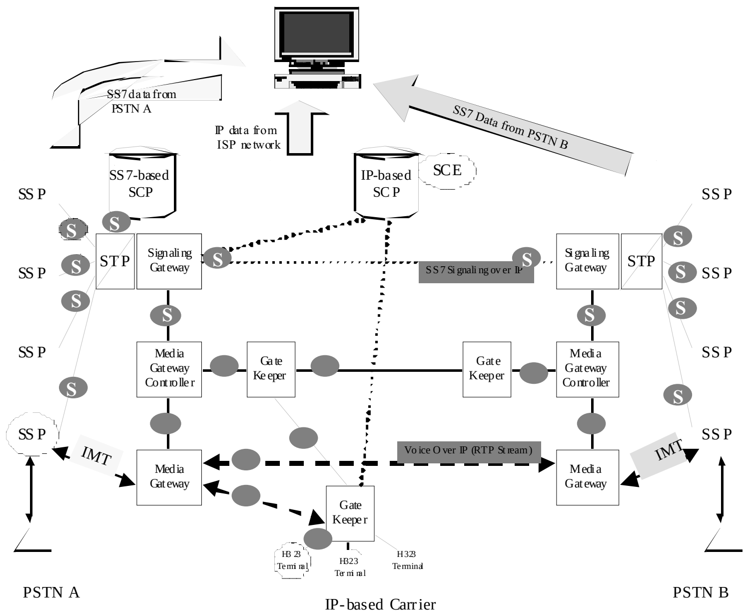

tain, and ensure the reliability of innovative telecommunication products and services implemented across SS7–based public switched telephone networks (PSTN) and next-generation wireline and wireless networks.

Operators use the monitoring device to correct network-wide problems from a central location and monitor and trace calls end-to-end. It enables a number of applications, including those deployed for fraud detection, mass call detection, billing verification, and quality of service (QoS) measurements.

When looking for such a system, operators should consider a device that provides complete monitoring capabilities for converged networks. Such a device should have a central, unified system to gauge both high-speed and low-speed links to generate data.

Operator should also look for features that have the capability to do the following:

- Monitor SS7 signaling and other telecommunications-related protocols transported over SS7 and IP networks

FIGURE 2

Telecom Service Surveillance and Maintenance System

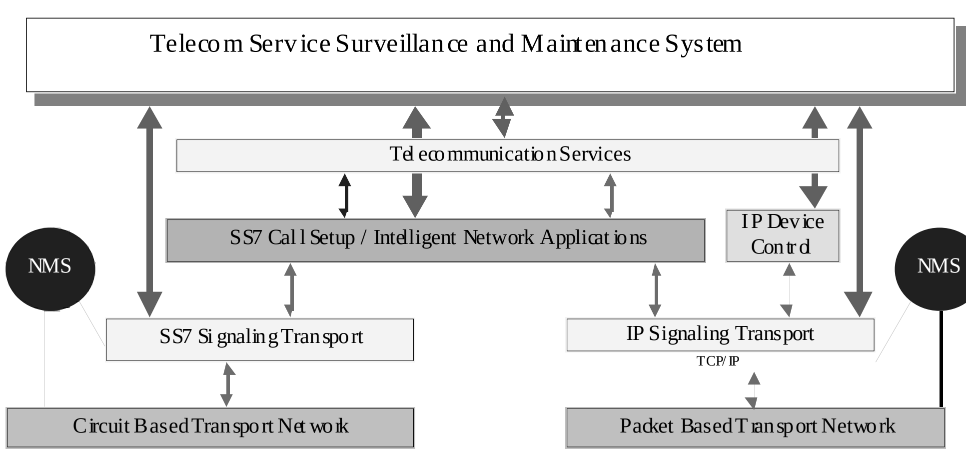

(A protocol-independent architecture allows providers to quickly adapt to new or modified IP–telephony signaling transport protocols.)
- Survey IP signaling transport protocols for faults and abnormalities
- Provide seamless protocol analysis across both SS7 and IP networks, assisting providers in maintaining and troubleshooting service-related problems.
- Perform proactive surveillance by troubleshooting problems and utilizing intrusive functionality across SS7 and IP networks
- Provide statistical traffic reporting for detailed traffic analysis across both SS7 and IP networks

In a perfect world, operators will be able to monitor, provide intrusive test generation, and complete protocol analysis, including filtering and call-trace capability. Additionally, a carrier would be able to gauge the performance of multiple protocols throughout its network, including TALI, MGCP, and SIP. Ideally, a single protocol analysis form allows the user to view messages and troubleshoot problems across the entire network, regardless of the interworking issues among various protocols. Real-time call-tracing capability allows the user to easily trace SS7 or personal communications services (PCS) call procedures. The user is able to trace calls involving multiple protocols, such as transactional capabilities application part (TCAP) interworking with integrated services digital network user part (ISUP), merely by providing the directory number of the call to be traced.

Looking Ahead

Future network operator developments include the incorporation of additional IP protocols, including stream control transmission protocol (SCTP), message transfer part (MTP)–2 user adaptation layer (M2UA), M3UA, MGCP, and SIP into networks. These protocols, like current IP and SS7 protocols, will require operators to dynamically shift their test and surveillance systems to meet these new challenges.

Increasingly, test and monitoring devices will rely on software platforms, which allow operators to purchase upgrades and add capabilities on an as-needed basis. In addition, current business applications, such as fraud detection and billing verification, are being expanded to incorporate data derived from emerging technologies and protocols. In today's quickly evolving networks, operators cannot afford to rely on subscribers to identify network problems. Via the deployment of a network-monitoring system that provides a consolidated view of a converged network, operators will be armed with the information they need to proactively manage their networks, services, and customer bases.

Non-Voice Network Management

Colin Pons

President

InJoy Communication BV

Abstract

The ever-increasing demand of new service offerings coupled with the ever-heightening cost-effectiveness programs result in an increased use of the installed intelligent network (IN). Cost-effectiveness often does not allow new service control points (SCPs) to be introduced. The result is that the operational risks are higher. These risks include not only failure, but also the size of the impact.

This paper discusses how congestion or overload conditions in the IN can be controlled. Several control methods are discussed, and a proof of the concept based on operational experience is presented. The objective of the presented concept is to assure stable performance of the network and a high quality of service (QoS).

Introduction

Like any other communications network, the IN is sized to carry the normal busy-hour traffic. A network such as this involves network elements (such as service switching points [SSPs], signal transfer points [STPs], and SCPs) and a communication interface (such as signaling system 7 [SS7]). At each of these four, abnormal conditions may occur that are not accounted for when designing the network. Or, it may occur for any reason that the number of queries is above the normal busy-hour traffic levels. For example, one can examine the SCP.

An SCP can handle only a finite number of queries (call attempts) from the network (SSPs). An IN should ideally be sized to allow the normal busy-hour traffic to be handled, taking into account at least a single network failure. If the number of queries exceeds the response capacity of an SCP, or a specific service on an SCP, the response either will be lost or will arrive too late to be handled by the SSP. This typically results in an error at the SSP network level and almost inevitably the calling-party redialing. This would cause a QoS degradation to be experienced by the customers.

Also, a SSP or a STP can handle only a limited number of queries. When the offered queries exceed this limit, it almost inevitably will cause calling-party redialing. These redial actions often worsen the congestion.

It is therefore necessary to implement sufficient measures to maintain a reliable service.

Intelligent Network Architecture

A typical IN may look like the one depictured in *Figure 1*.[1]

Each of the three nodes and the communication interface are shown in *Figure 1*. Typically, the SSP co-resides with the CCF on an exchange. This can be an exchange of any type. The STP can perform several roles. Among these are the traffic-load distribution, fall-over capabilities, and query routing (to the appropriate service).

The SS7 network provides for a standardized interface between these network functionalities. SS7 also provides the routing capabilities between nodes. Two types of routing are applicable within this network: routing on SSN, which uses the physical point code, or routing on GT, which uses the Global Title (a logical node identifier) for routing.

With respect to the SSF, losing responses from the SCP or responses from the SCP arriving too late is not an efficient way of handling an overload situation. An SSP will not release its resources until a timer used to protect against lost responses, known as T_{SSF}, has expired. The T_{SSF} timer is typically in the order of seconds, whereas the normal response time from the SCP is in the order of tens of milliseconds. Therefore, significant SSP resources can be wasted in the case of an SCP failing to respond. The outcome could be that the SSP will not have adequate resources to trigger an alternative SCP or, worse, may become overloaded and fail to handle other calls, whether IN or not. Similar concerns can be raised for the STP depending which functionality is used.

The SS7 intelligent network application part (INAP) stack utilizes the signaling connection control part (SCCP) layer for routing messages across this network. The SS7 links and SCCP routing tables are ideally suited for normal busy-hour traffic.

An SCP (or specific service) is considered overloaded when it cannot respond to a service query within a reasonable time (typically less than 100 ms). An SCP may become overloaded for a variety of reasons. For instance, if an SCP, or specific IN service, cannot process queries at a rate faster than they are being sent, then the SCP, or the specific service, is considered overloaded.

In this network model, a number of bottlenecks can be identified:

FIGURE 1

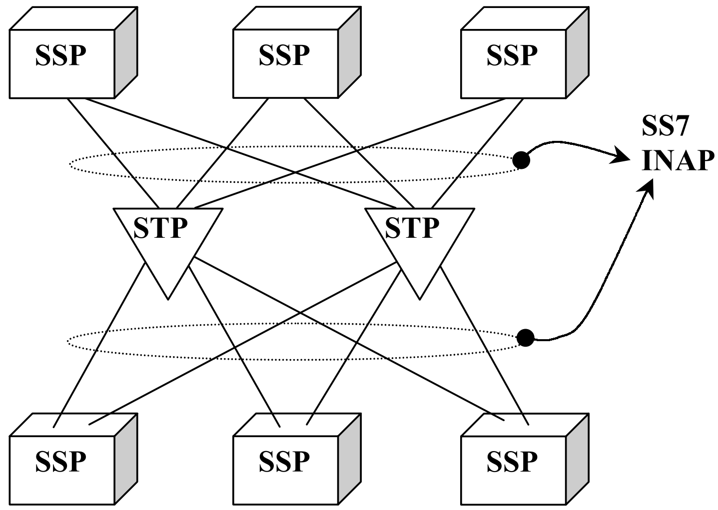

- Software resource shortage due to a sudden increase in queries possibly leading to one overloaded service affecting another and/or SSPs becoming overloaded
- An increase in central processing unit (CPU) load due to a concentration of events or a rapid build up of queries; SCP becoming unstable, possibly leading to failure
- A shortage in transmission capacity; the SS7 network becoming congested
- Routing errors, which are man-made errors

To solve the above, the following should be done:

- A reduction of the number of queries should be effectuated as close as possible to the source. This would help to eliminate unnecessary resource usage, allowing more successful calls to proceed.
- The effect of a disruption or a sudden increase in the number of queries for a service should be confined. The effect should be prevented from interfering with other elements of the network or other services.
- Any abnormal trend in query volume should be detected as soon as possible to allow early intervention.
- Better network resource usage efficiency.

Current Measures

IN Call-Gap Operation

IN call gapping is a mechanism to prevent SCPs (or specific IN services) from becoming overloaded. When IN call gapping is active, excessive queries, which would normally be sent from the SSP to the SCP, are throttled at the SSP by enforcing a minimum time interval between queries being passed to the SCP for processing, thus allowing the SCPs to process queries at its maximum supported rate. Calls gapped at the SSP[2] will receive a defined treatment known as the gapTreatment—e.g., an announcement, or more likely a tone, is played to the calling party. The SCP tells the SSP what to gap for in the gapCriteria (see ITU Recommendations Q.120x for more details).

IN call gapping is *not* the same as public switched telephone network (PSTN) call gapping and does not prevent SSPs from becoming overloaded. It is purely an SCP and SCP service protection mechanism.

IN call gapping can be utilized in two ways: manual call gapping and automatic call gapping.

Manual Call Gapping

This form of call gapping is instigated by operational personnel following understood procedures for a number of *anticipated* events. For example, the traffic-management group may anticipate an excessive call volume due to television quiz shows or special promotions taking place for particular service numbers. In this case, manual call gapping is set up ahead of the event to prevent SCP overload. Manual call gapping may also be established to protect SCPs during planned network outages or upgrades, when the overall available IN capacity may be reduced.

Automatic Call Gapping

This form of call gapping is instigated automatically from the SCP network in abnormal or unpredicted situations. Once an automatic call gapping event has occurred, the cause of the overload should be investigated and manual call gapping procedures put in place if required. Calls gapped at the SSP will receive a defined treatment—the gapTreatment.

Notice in *Figure 2* that immediately after the overload condition subsides (i.e. the number of ISUP IAMs drops to a normal level) the call gapping is still active. Therefore if any calls arrive within the gapInterval they will still be gapped. After the duration time has expired the SSP removes the call gap and no further calls will be gapped.

SS7 Fail-Over Capabilities

It is assumed that the reader is sufficiently aware of the capabilities of SS7. SS7 allows a message transfer part (MTP)–level routing of messages over several routes and

FIGURE 2

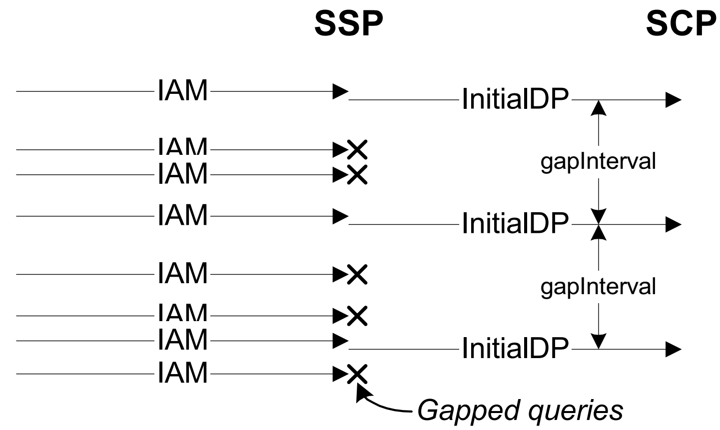

link sets in a redundancy scheme. If one route or a link set fails or is congested, another route could be selected. At the SCCP level, a similar redundancy scheme is implemented. In the case of routing on GT, a secondary or back-up address can be used for redundancy.[3] At the MTP level or the SCCP level, there are no dynamic routing capabilities. The settings are operator defined and can only be changed by an operator or from outside the SS7 network. SS7 provides very reliable communication but is limited in network resource utilization.

Congestion Control

The IN consists of the SSP, STP, and SCP. Interestingly, the SSP and the STP are often special-purpose computers and telephony exchanges, whereas the SCP could also be implemented on a general-purpose computer. Actually, virtually all newly implemented INs use general-purpose computers for the SCP. Furthermore, it is important to notice that there are far more SSPs deployed than SCPs. In fact, at the SCP, an accumulation of queries from a number of SSPs takes place, and SCPs are more likely to first experience abnormal query volumes than any other element in the IN. Given these reasons, it is appropriate to take the SCP as a target node.

Detecting an Overload Situation

The SCP could be seen as a processing system that is built up out of a buffer for storage of queries that are awaiting processing, a mechanism for processing these queries (SLPI), and a CPU for running the service logic (SLPI) (see *Figure 3*). The number of messages (MSUs) waiting in the service logic program instance (SLPI) input and output queues can be used to determine if a service is overloaded or not. In practical terms, only the contents of the input queue need to be considered.

The queues or buffers reflect all buffers present in the system, such as the input/output (IO) buffer (SS7 link), the dispatcher buffer, the database queue, etc. The input queue reflects the number of queries waiting to be processed by a SLPI. The number of pending items in the queue will vary depending upon how quickly the SLPI can service the requests. If the rate of queries is greater than the rate of service by the SLPI, then the number of pending queries will increase. The SLPI will request the CPU for a time slice each time it wants to process a query.

Using queue size to determine overload conditions means that virtually all other external factors are automatically taken in to account. If the number of call attempts is normal but an abnormal event takes place on the SCP that reduces the amount of CPU available for a service, then the number of pending queries in the queue will grow.

Controlling an Overload Situation

As mentioned earlier, traffic congestion and network congestion can occur in the following instances:

- A software resource shortage due to a sudden increase in queries
- An increase in CPU load due to a concentration of events or a rapidly built up of queries
- A shortage in transmission capacity, eventually reflected also by a software resource shortage

As discussed in the preceding section, a software resource shortage is reflected by the input and output queue occupancy. Either the SLPI cannot handle all incoming requests or the SCP cannot transmit all messages. The latter reflects a shortage in transmission capacity in the downlink to the SSP. There are several reasons that contribute to the ability of a SLPI to service the outstanding requests. The main factor, however, is the number of threads that can be supported and access to the CPU. This requires monitoring of the queues and control of the flow of queries.

CPU occupancy increase causes processing delays. This will harm SLPIs trying to gain access to the CPU. Eventually, CPU overload conditions result in an increase in uncompleted calls. To prevent an unhealthy concentration of queries, it is necessary to determine the maximum number of queries that can be accepted based on CPU load and where the SCP remains stable. This requires control of the number of queries toward a SCP.

FIGURE 3

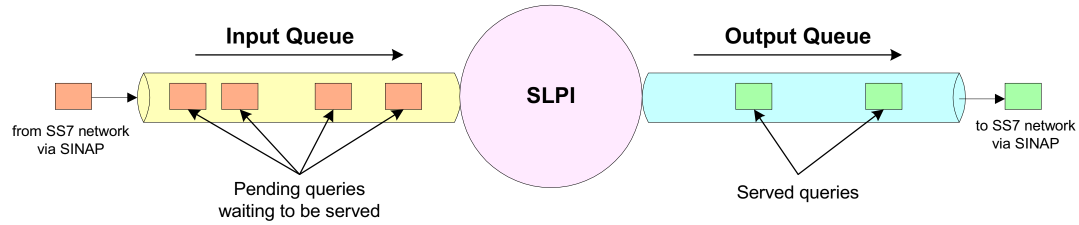

From the conclusions for both reasons for congestion, it can be concluded that the flow of queries needs to be controlled. IN call gapping limits the number of queries and is suitable for preventing software resource shortages that cause congestion of the network. But IN call gapping limits the number of queries from an SSP toward an SCP. The problem with this is that other SCPs may well have the available software resources to satisfy the offered load from that SSP. IN call gapping is not suitable for limiting the *total* number of queries from multiple SSPs toward an individual SCP.

The IN call-gap operation needs to be backed up by a flow control mechanism.

Flow Control

The flow control manages the number of queries and their routing through the network. In contrast, the IN call-gap operation only limits the rate with which an SSP sends queries to an SCP. The flow control would allow for a better utilization of network resources and higher successful completion rates for calls.

Criteria for Control

Tests and network statistics have shown that the queue size starts to increase as the CPU load reaches 75 percent to 90 percent, depending on the circumstances such as high provisioning rate. A general-purpose computer is best operated not higher than 85 percent to allow necessary system processes to have access to the CPU. To allow a margin for handling an increase in queries (including redundancy) as well as provisioning processes and system management processes to take place, it is recommended that an SCP's CPU load is no more than 60 percent in normal busy-hour operation. Therefore, the queue size has a direct relationship with CPU occupancy.

In the IN, this means monitoring and use for flow control of the following:

- The offered queries per SSP
- The received queries per SCP (and if possible per service or SLPI)
- The CPU load
- The SS7 link occupancy
- If possible, the queue length of each SLPI at each SCP (system dependant)

Implementation

Currently, most networks already have network-management systems that collect the aforementioned data. If this is not already done, this information may be obtained from the network elements themselves. It is important to notice that the statistics should be readily available. A time lapse of 15 minutes would prove unsatisfactory in real life. Events that reshape traffic loads and possibly effect network QoS can take just a few minutes or even less. It is practical to have new statistics available once a minute. This would satisfy the need to have real-time data available and not to stress the network elements with obtaining statistics.

One other important issue to notice is that the flow control should be done outside any SSP, SCP, or STP node with the network. Each of them might fail, possibly impacting the flow control adversely. The flow control should be implemented at network-management layer as a separate system.

The flow-control management system could possibly interface with the network through the SCP or the STP, depending the exact network configuration. This interface would allow the initiation of IN call-gap operations. Another type of interface is necessary to manipulate SCCP routing tables at SSP, STP, and SCP.

Control Application

In a first implementation, the non-voice traffic-management system (see *Figure 4*) comprises two parts: one part for obtaining all statistical and SCCP routing information and sending SCCP routing settings and another for simulation. In the simulation, the effect of failing SS7 links, STPs, and SCPs can be examined. The simulation results can than be used for effective network control. The system is fed with information from the SCP (via the SMS) and from a link-monitoring system. The link-monitoring system allows the determination of the sending SSP's identity. The routing tables present in the different nodes are manually fed to the system. The results so far indicate that a good prediction of network behaviour can be made when the different elements fail. There is also a clear indication of what can be done to minimize the effect of multiple failures in the network, such as the initiation of an IN call-gap operation or changes to the SCCP routing tables. The first implementation is to serve as a proof of concept. The objective is to have the real-time part determine critical situations and to have it calculate the appropriate actions.

There is some issue with dynamically changing the SCCP routing tables. From an operator point of view, the operator is no longer aware of the actual settings each time. Furthermore, dynamically changing routing tables may have an adverse effect. We have avoided that by properly select-

FIGURE 4

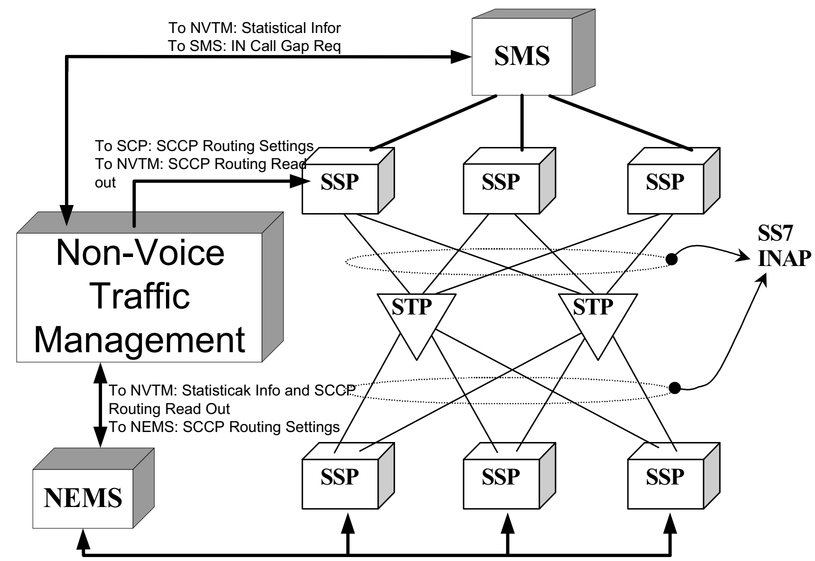

ing the algorithm and by implementing thresholds and limiting the number routing table changes per time period.

Conclusions

This paper has defined the most prevailing causes of congestion and QoS deterioration in an IN. The uses of IN call gapping and the flow-control concept, and their implementation have been described. Flow control allows for a better network utilization and better network control in cases of failures. In the future, it will be necessary to develop an automated non-voice traffic-management system. The algorithms are available, but they need to be implemented and fine-tuned to provide the desired results.

The ever-increasing demand of new service offerings and the cost-effectiveness programs in place at virtually all telecom operators require that current investments be better utilized. This is only possible if the appropriate measures are taken to prevent a failure from having a detrimental impact on the experienced service offering. We believe that a non-voice traffic-management system plays a key part in ensuring QoS despite increased demands on the IN.

Notes

1. For reasons of simplicity, elements such as special resource functions (SRFs) are omitted. However, this discussion also applies to them, since the same issues are involved. The conclusions can also be applied, therefore, to IN nodes such as the SRF.
2. Not all SSPs interpret the IN call-gap operation, and the *most specific match* rules, in the same manner. Therefore, it is important to understand the call-gapping behavior of the specific SSP before gapCriteria can be applied.
3. If the sending node is aware that the next node is the end node (SEP), the recommendation is to use routing on SSN. In this case, the end node's point code is used to address it. There is no SCCP back-up route applicable in this scenario (although one could be provided through the MTP level). If the sending node is unaware of the status of the next node, or if it is aware that the next node is not the end point, the recommendation is to use routing on GT.

OSS Flexibility and Interoperability

Richard G. Ramos
B/OSS Solution Architect
EDS

Overview

Few telecommunications and Internet service providers (ISPs) are immune to the unrelenting pressure to improve their earnings performance, leverage their existing network infrastructure, improve customer satisfaction, and bring new services to market more quickly. They endure these pressures—all while striving to reduce their capital and operating costs in a chaotic, unpredictable market.

To meet these challenges, next-generation service providers are realizing great value by deploying a new breed of operations support systems (OSSs) based on service-enabling platforms (SEPs).

SEPs are Web-based middle-tier software and hardware technologies that rapidly deploy and change new service models. They instantly deliver dynamic, personalized bundles of enhanced-quality voice, video, data, and content services over any browser-based appliance. SEPs can provide unprecedented levels of scalability, flexibility, and repeatability by applying e-business and Internet technology to solve some long-standing problems that plague next-generation service delivery.

This white paper examines the application of Internet infrastructure technology to OSS deployment and addresses two of the most expensive and painful problems with next-generation OSS implementation: multivendor interoperability and flexibility.

How the Internet Has Transformed OSS

Before the Internet, OSSs managed circuit-switched voice services and their networks in a tightly regulated, anti-competitive environment. The resulting systems became increasingly monolithic environments in functional silos of the incumbent carriers.

In the pre-Internet days, each department—whether customer service, sales, billing, provisioning, or order entry—defined its OSSs according to its organization and within the scope of specific services. In effect, wireless, high-speed data, voice, business, and consumer services all resided within different organizations. Each service was often its own profit-and-loss center with its own information technology (IT) infrastructure and budget. OSSs evolved as point solutions based on fixed, controllable, predictable business models.

The Internet and e-business revolution, however, has shattered any hope of predictable business models. The new digital economy demands telecom e-business services that require collaboration between trading partners (such as application service providers [ASPs] and content providers) to deliver the next-generation services that their customers need.

Service providers competing in the digital economy must endure chaos, uncertainty, intense competition, and rapidly evolving, diverse technology options. Two of the most expensive and painful problems facing service providers' OSSs are multivendor interoperability issues and system flexibility. In fact, these problems are gateway factors in the successful deployment of next-generation services.

The Importance of Interoperability

Achieving interoperability, or system integration of multiple OSS components, is a problem rooted in the evolutionary nature of OSS technology. The functional silos have often made it necessary to integrate 10 or more systems for the basis-fulfillment, assurance, and billing functions. The following are some of the problems inherent in achieving interoperability:

- ***Integration:*** Costs are typically $1million x (N x [N-1]), where N = the number of systems that must be integrated[1]. This expense can result in first-time OSS project implementation costs as much as three to five times the system's software license cost and annual operating costs as high as the software license costs.
- ***Implementation:*** Time frames may stretch into years. Implementing a provisioning system in a legacy environment today means that the service provider must migrate product catalogues and customer databases to a new environment—and in some cases several environments. This data migration often adds many months to an integration project.
- ***Order-to-Fulfillment Completion Rates:*** Order completion rates are fundamental to the efficiency of a service provider's operations. Today's service providers are doing well if three out of 10 orders make it successfully to completion. Effectively coordinating an order across multiple systems—such as order entry, order management, billing, provisioning, activation, inventory, and sales-force automation—depends on how well these functions are integrated. Any compro-

mise affects the time and cost of successfully fulfilling an order.
- ***Data Synchronization and Maintenance:*** The more systems that must be integrated, the greater the likelihood that multiple databases must be maintained. If the systems require redundancy, OSS operating costs increase significantly. Add these maintenance costs with software revisions of multiple systems and upgrades, and the total cost of OSS ownership increases to 75 percent of the initial capital expense (CAPEX).
- ***Service Fault Root-Cause Determination:*** Correlating a service or network fault to an affected customer—and then calculating and applying appropriate credits to that account (based on a service-level agreement [SLA])—is a primary objective of service-assurance flow-through. Sixty percent of all time spent dealing with any given network fault is spent pinpointing its root cause. That means 60 percent of the network's fault-related downtime can be attributed to root-cause analysis not happening fast enough. Faster, more automated root-cause analysis can reduce network downtime by half—a significant competitive advantage in a market that puts a premium on SLA guarantees.[2]
- ***Comprehensive Customer View:*** A comprehensive customer view requires real-time access and synthesis of all pertinent customer information, regardless of where it resides. The degree to which all service-provider entities have a consistent, thorough, and current view of the customer and the resources required to serve that customer determines the level of customer service and satisfaction that providers can deliver. By taking a 360-degree view of customers, service providers can market to them on a one-to-one basis, strengthening the relationships, improving revenues per customer, responding to their problems satisfactorily and, ultimately, ensuring their long-term loyalty.

Many of these interoperability problems have their genesis in the attitude OSS vendors take toward interoperability, which is often regarded as someone else's problem: "We provide an application programming interface (API) the customer integrates to, and it is their problem to make sure it works."

Today, most service providers still choose to procure tactical point solution OSSs according to function and price, and at the expense of the entire integrated solution. Little wonder that few OSS implementations demonstrate the value expected of them at the outset. Clearly, the point solution versus an entire integrated approach has not served the industry well.

Fortunately, OSS knowledge is growing. Service providers now have the tools to understand and measure the impact of interoperability problems and to transform them into opportunities such as flexibility.

The Importance of Operational Flexibility

Flexibility benefits the bottom line by increasing customer loyalty and reducing operating costs, total cost of ownership (TCO), and time to profitability for new services. Flexibility is vital to a just-in-time service provider, particularly one in which the OSS enables rather than impedes the business. In the chaotic, unpredictable environment of a collaborative digital economy, the next "killer app" is unknown, but service providers who adjust rapidly when they see the killer model starting to emerge will be the ones who rapidly gain market-share.

Flexibility problems not only increase operating costs related to system changes but also bottleneck time to market for new services. These factors increase lost-opportunity costs and frustrate IT managers in their efforts to more effectively serve the customer.

To improve earnings performance, next-generation service providers must increasingly and continuously differentiate themselves by doing the following:

- Offering premium, dynamic, personalized bundles of voice, video, applications, and content services
- Transacting over the Web (configured, procured, paid, activated, supported, and assured)
- Deploying over any type of infrastructure (circuit- or packet-switch, optical, wireless, cable, satellite, or copper)
- Functioning across a complex value chain
- Ensuring five-nine (99.999 percent) reliability and quality just as traditional voice services do

Services that differentiate one company from another include secure, guaranteed quality of service (QoS) virtual private networks (VPNs) for enterprises, on-line gaming, instant messaging (IM), and streaming video/audio for consumers.

Because the benefits are often difficult to measure and quantify, they tend *not* to be used as the primary purchase criteria for an OSS solution. New economic models, however, have begun to reflect these benefits in faster time to profitability, accelerated earnings, higher revenues per customer, and reduced customer churn. It is vital that service providers consider the importance of flexibility in their OSS evaluations.

Consolidation Makes Interoperability and Flexibility More Critical

Interoperability and flexibility will become even more vital in the coming wave of consolidation among telecommunications and ISPs. Carriers looking to offer new services or enter new markets are opting to merge, acquire, or partner rather than grow organically. Such options may speed market entry, but they wreak havoc on business support systems (BSSs) and OSSs, and on the ability to adapt and change in a chaotic world. Network integration is difficult enough with an array of devices such as 4ESS, 5ESS, optical, edge/core routers, gateways, softswitch and media servers. Further integration at the service and business-systems layer, including multiple billing systems, order management, provisioning, activation and assurance—and all for different services and physical networks—magnifies the complexity. One global service provider that followed this integration strategy now faces more than 50 different billing systems alone!

Internet Technology Applied to OSS

The Internet and e-business revolution has spawned a new generation of enterprise operations technology. Innovations such as enterprise application integration (EAI) middleware, application-server technology, Java, extensible markup language (XML), and a host of other tools offer a new alternative to OSS implementation and help next-generation service providers meet the demands of their business and consumer customers.

Web server, J2EE, Java, XML, enterprise workflow, and rules-based environments as the building blocks of e-business are ideal platforms for OSSs in the collaborative digital economy. BEA, IBM, HP, iPlanet, Oracle, Sybase, and Fujitsu have each introduced their own version of these platforms, and industry experts already regard these platforms as revolutionizing enterprise computing. When these proprietary platforms are applied to OSS implementation, however, their effect is *transformational* and *disruptive*.

A new breed of OSS platforms applies these e-business components to solving the next-generation OSS challenges of zero-touch, end-to-end, flow-through provisioning; single-point-of-entry; Web-based self-care; guaranteed QoS; value-based pricing; and dynamic, personalized bundles of voice, video, applications, and content services. These new OSS platforms meet the demands of today's services and, more important, are flexible enough to address services yet undefined. We call these service-enabling platforms.

SEPs apply the TeleManagement Forum's e-business Telecom Operations Map (eTOM model) to address the business issues facing next-generation service providers. These platforms couple the network and service (the operation) with the business layer (the customer) by correlating customer-affecting issues with network events.

SEPs are designed with service order-to-fulfillment at their core. Generic business processes, rules, and workflow templates for specific types of providers are available out of the box to dramatically speed implementation. Workflow or rules can be added or modified according to specific business requirements, which speeds implementation and reduces recurring and non-recurring costs.

More important, SEPs shorten the service order-to-fulfillment lifecycle by enabling the interoperability of all customer-service resources to a specific customer's order request. In the telecommunications industry, fulfillment response time and the market demand for products correlate directly. The faster you can activate customers, the higher the service volume per customer.

Another subtle but important benefit of SEPs is their ability to provide accurate enterprise health checks through a business-intelligence or enterprise-analytical application. The rich information these platforms provide allows managers to examine all levels of the enterprise and understand root causes of service and performance problems, such as customer dissatisfaction and negative earnings. This deep, timely information speeds remedial and proactive measures to improve enterprise performance and customer satisfaction.

The Components of Service-Enabling Platforms

SEPs take an enterprise-wide approach to solving the OSS flexibility and interoperability problems plaguing the industry. A SEP involves the following e-business platform components and their relationships (see *Figure 1*).

Enterprise-Wide Process Automation

Enterprise-wide process automation allows interoperability between front- and back-office applications at multiple decision levels. It also instills the ability to instantly modify

FIGURE 1

Enterprise-Wide Automated Processes

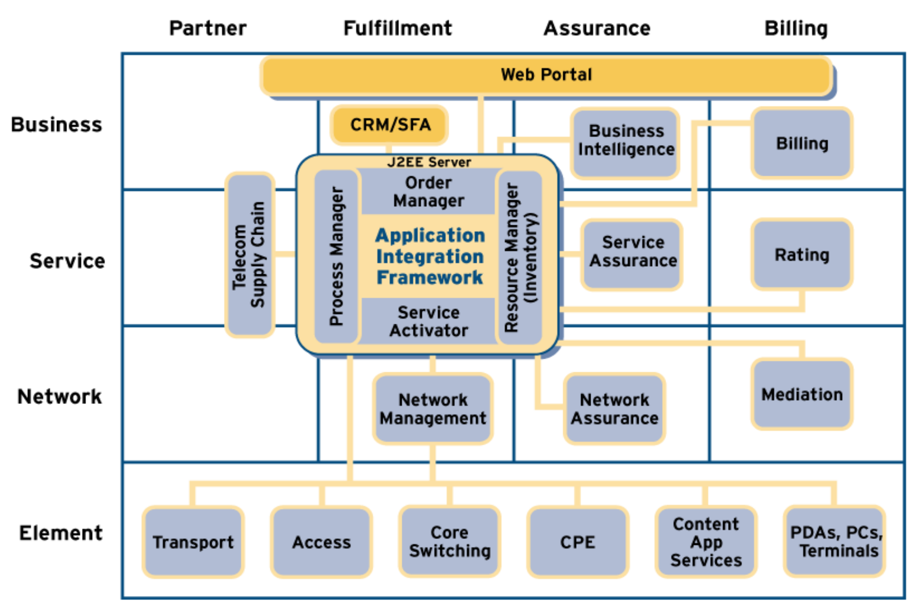

workflow and business rules at the enterprise level, which is vital for rapid response to market and business changes. Enterprise-wide process automation requires that all applications and data be controlled directly by the business rules and workflow engine (process manager). Applying and enforcing rules and workflow to enterprise business processes can reliably invoke other activities—such as automatically sending confirmation e-mails to customers, notifying account managers to contact customers, and so on—that can ensure consistent, superior service quality, fewer operational errors, and enterprise-wide control across a broad organization and supply chain.

Application Integration Framework

The application integration framework accepts data from legacy or other OSS applications in native format. By accepting data in native format using plug-in adapters developed in XML, Java, or Corba, massive data migration is not necessary before the SEP can deliver service. The plug-ins are used only to expose the resource object to workflow and contain no embedded logic. The result: significant reduction in the time required for integration with proprietary application programming interfaces (APIs). Removing data migration from a legacy OSS interoperability project reduces implementation time frames by months.

J2EE Application Servers

J2EE application servers handle many of the functions that were traditionally coded in the OSS application's proprietary framework. They also handle the functions that were managed using traditional message brokers. Load balancing, persistence, state machine, directory and security services, messaging, transactions, and database access are out-of-the-box J2EE services. J2EE application servers support XML, Java, extensible server page (XSP), extensible stylesheet language transformation (XSLT), and Enterprise Java Bean (EJB). These J2EE servers, in combination with the supported applications, simplify OSS development and integration. While traditional OSS vendors have developed such functions internally, leaving the customer with the integration challenge, J2EE application servers form the core technology upon which the next-generation OSS can be built.

Document Object Model

The document object model (DOM) is an EJB supported in the J2EE application server. The DOM dynamically transforms data into a format that can be used to synchronize disparate data sources in a complex OSS environment. Data synchronization is vital for enterprises with multiple OSSs that handle overlapping functions such as billing. The DOM makes dynamic data views possible by transforming data from disparate OSSs and using XSLT and hypertext markup language (HTML) to present a common look and feel to users.

Web-Based versus Web-Enabled

Next-generation OSSs developed with J2EE application servers and Web servers provide an out-of-the-box interface for accessing data from a variety of appliances. The devices merely need connectivity and browser capability. All users receive the same graphical user interface (GUI), regardless of the device.

Core OSS Components (Order Management, Service Inventory, and Service Activation)

Traditional OSS implementations integrate disparate service-fulfillment components such as order management,

FIGURE 2

Service-Enabling Platform Implementation Methodology Comparison

Conventional OSS Implementation Model	*Service-Enabling Platform Implementation Model*
System design 80% completed up front	System design 20% completed up front
Heavy up-front business process engineering	• Out-of-the-box embedded business processes • Iterative process engineering and refinement
Heavy front-end requirements analysis process	Rapid prototyping via multiple configuration levels: o Business processes o Inventory object modeling o Presentation layer o API plug-ins
Best-effort (no guarantee) integration to multiple APIs of adjunct systems	Well-understood integration framework that ensures end-to-end business performance is not compromised through integration
Expensive change-management process for specification changes	Changes expected in the implementation process

service provisioning, inventory management, and service activation. The decoupling of these core components is a significant cause of OSS integration problems. SEPs allow for a tighter coupling of these functions using enterprise-level business rules and workflow. In comparison to traditional approaches, SEPs turn issues of multiple product catalogue aggregation, data migration, and application interoperability into minor inconveniences.

Implementing Service-Enabling Platforms

As is true with all enterprise IT, achieving successful results by adopting SEPs is as much a function of their implementation as of technology itself. In this new economy, OSSs can no longer be designed or implemented using traditional best practices, which assume a stable service model. Service models can change overnight, and even long-term models are difficult to predict[3]. OSSs have to keep up with the business, and keeping up requires methodologies that facilitate rather than impede implementation and that minimize one-off customization. To reap the full benefits of easy interoperability and rapid flexibility, SEPs require new implementation models. *Figure 2* compares these new methods with conventional ones.

A Look at the EDS B/OSS Solution

EDS' BSS/OSS (B/OSS) Application Interoperability Framework (AIF) is a Web-based OSS that utilizes a SEP. EDS, in conjunction with several solution partners, has developed a fully interoperable Web-based OSS for integrated communications providers (ICPs). The B/OSS is a highly scalable and robust OSS solution that EDS provides as a hosted service, charging clients on a pay-as-you-grow model. The core B/OSS modules are Order Management, Service Inventory, and Service Activation. These modules integrate with the J2EE application server through the AIF. Application interoperability is choreographed by the Process Manager (business rules and workflow engine). EDS worked with solution partner—Digital Fairway—to create Process Manager templates that decrease time and expense for implementation.

Conclusion

Next-generation service providers rely on OSS vendor support to meet the changing business climate demanded by our collaborative digital economy. As deregulation and competition continue to erode transport prices, providers must rely on value-added services, collaborative trading partners, QoS, speed to market for new services, and a quality customer experience as market differentiators.

SEPs at the core of providers' OSSs will result in significantly improved earnings performance by reducing integration costs, speeding new service time to market, improving customer satisfaction and loyalty, and increasing revenue per customer.

Notes

1. EDS Service-Enabling Platform Economic Model, Vo1. 8
2. James Montgomery Scott, "Service Assurance in an Era of Uncertainty," June 28, 2001.
3. Remember the integrated services digital network (ISDN)—the "best-planned failure" in telecom—when you think of third-generation (3G) wireless. Remember also the Internet—the "worst-planned success" in the telecommunications industry.

Design Imperatives for Survivable Networked Systems

Ramana Reddy
Director, Concurrent Engineering Research Center, Lane Department of Computer Science and Electrical Engineering
West Virginia University

Sumitra M. Reddy
Research Associate Professor, Lane Department of Computer Science and Electrical Engineering
West Virginia University

Srinivas Kankanahalli
Research Associate Professor, Lane Department of Computer Science and Electrical Engineering
West Virginia University

Introduction

Rapid proliferation of networking technologies, ubiquitous computing nodes, dispersed storage services, and advances in markup languages are contributing to the transformation of conventional enterprises into digital enterprises. These digital enterprises should support high availability, integrity, and confidentiality of their digital assets. As this transformation continues, protecting the digital assets of an enterprise from compromising external/internal attacks and ensuring continued operation in the presence of partial degradation of the system assume paramount importance. While a networked system is inherently more capable of continued operation and graceful degradation in the presence of faults than a centralized system, it is also more vulnerable to malicious attacks, since there are many more points of entry. Furthermore, the inherent distributed nature of large-scale public networks such as the Internet prohibits any central enforcement of any single security policy. However, any attack on the digital assets of an enterprise results in a disruption of services within as well as outside the enterprise, resulting in a significant loss in productivity and revenue. A designer of a digital enterprise system is faced with balancing these countervailing forces.

What Is a Survivable System?

We begin by observing that a networked digital enterprise system exists in a hostile environment where every transaction is presumed to be "suspect," whether originated internally or externally. A system may be considered to be survivable if it can continue to function as defined in the presence of malicious attacks, failures, and accidents by incorporating redundancy, recoverability, and intelligent management of transactions. In other words, a survivable system guarantees that the digital assets are highly available and highly secure, maintains integrity, and ensures system confidentiality in the face of malicious attacks. In the remainder of this paper, we briefly describe a number of design and deployment imperatives for survivable networked systems:

1. Redundancy
2. Recoverability
3. Quarantined Transaction Execution
4. Transactions and User Profiling
5. Virtual Node Isolation
6. Audit Nodes
7. Dual-Track Execution
8. Predictive Models
9. Signature Matching
10. Ephemeral Entry Nodes

A brief description of each of these design imperatives follows.

Redundancy

This is one of the foundations of reliable systems. Traditional engineering systems incorporate more than one component or subsystem for each function to deal with failures. They are active only in tandem—not simultaneously. However, networked digital systems must have redundant subsystems that are simultaneously active. These subsystems must be not only physically apart, but also designed in such a way that all of them are not susceptible to the same type of attacks. Examples of redundant technologies in an enterprise include the following:

- *Redundant Array of Independent Disks (RAID) Arrays:* These can support mirrored data and also have the ability to reconstruct data that is lost due to failed disks (RAID level 5).
- *Web Farms and Clustered Computing:* In most large organizations, there are Web server farms for serving high-volume Web transactions. Usually the server farm consists of many servers performing the same function managed through a load balancer. Cluster computers such as Beowulf Clusters consist of clusters of commodity processors connected by commodity networks to perform a complex task. These clusters can also be used as an enabling technology for survivable systems.
- *Clustered File Systems:* These provide applications with a unified view of a single system even though the actual data and the metadata are distributed across multiple nodes in a cluster. An example of a clustered file system is the Petal file system from Compaq Research.

Recoverability

In a digital information system, each transaction alters an arbitrary number of data objects. By the time a "rogue transaction" is discovered, it may have already altered a number of data objects that may be distributed across the enterprise. To ensure system integrity, the design must incorporate techniques for undoing all the "suspect" transactions and also establish a checkpoint from which the system can move forward with the assurance of system integrity. Some technologies that may be used include the following:

- *Journaled File Systems:* The basic idea of a journaled file system is to maintain a log of all "write" commands to a file system in a journal that is separate from the file system. Upon a system failure, the system can be checked for the status and restarted from an earlier safe point. Such log structures are routinely used in databases for recovery of transactions.
- *Snap Shots:* Here the metadata of a file system is frozen and accessed as it existed in some earlier point in time. A product such as the Network Appliance WAFL™ file system uses this approach for data filing and recovery.

Databases routinely provide checkpoints from which they can recover when a transaction aborts. Databases also use "shadow" transactions and "two-phase commit" protocols to support recoverability.

Quarantined Transaction Execution

Since our basic premise is that "all transactions are suspect," the design of a survivable system should incorporate techniques for the execution of transactions in "quarantine," i.e., in isolation from other parts of the system. The effects/results of these transactions may be migrated to the rest of the system only after validating that the set of transactions have not compromised or corrupted the system. This may be thought of as a system that moves in steps, while the transactions are executing continuously. The quarantined transactions may be executed by a separate "processor-memory-network" combination to obtain complete transaction isolation.

Transaction and User Profiling

Every enterprise could be described by the type, number, and distribution of a set of transactions that may form "patterns." As transactions are executed, there could be "watchdog" processes to monitor the activity. These monitors could interact with a knowledge-based system that analyzes the transaction patterns to detect any unusual activity. This could be interpreted as telltale signatures of undesirable or hostile transactions. In a like manner, we can also assign "patterns" to user types. When significant deviations from ad hoc patterns are detected, a compromise could be inferred. Lessons learned from security organizations could form the basis for these knowledge-based systems.

Virtual Node Isolation

A networked system by definition has a number of discrete nodes that are distributed across the network. Survivability of such a system can be enhanced by making it more difficult for an intruder to detect the related nodes of the system. Inter-node communication could be mediated by "black boxes" that encapsulate the origin/destination node mapping and are designed to resist probing. An attempted probe on a "black box" could automatically signal/infer the presence of hostile activity. To make the discovery of mappings more difficult, they could be made dynamically driven by secret algorithms.

Technologies such as virtual local-area networks (VLANs) in Ethernet or Zones in Fibre Channel technologies may be used as the basis for achieving virtual node isolation.

Audit Nodes

A survivable system should incorporate specially designed nodes whose sole purpose is to audit the changes that are taking place at performance nodes. When an audit node discovers compromised data, it will alert every performance node to shift into an enhanced security mode. These audit nodes should be designed in such a way that they cannot themselves be discovered and compromised. These audit nodes can also dynamically change roles to thwart the discovery of audit patterns.

Dual-Track Execution

In multimedia systems, the encoded data is often transmitted in parallel streams: one with high quality and another with a lower quality. When parts of high-quality data are lost, they can be substituted with corresponding lower-quality data so that the signal continuity is maintained, albeit, with a somewhat reduced quality. In a similar manner, under some circumstances, we can incorporate dual execution of transactions that are independent of each other but can be mutually substituted in order to repair damage caused by hostile acts.

Predictive Models

A simulation model of a distributed information system at multiple levels of abstraction may be created and executed in parallel with the real system being tracked. At any point in time, key performance attributes of the simulation model could be compared with their counterparts in the real system to ascertain unusual deviations or compromised states. These models could serve as external triggers to audit nodes to initiate consistency checks. Some of the intrusion detection systems use a similar technique to alert statistically significant behavior.

Signature Matching

Many attacks on information systems follow "predictable" patterns. By analyzing the transaction patterns, the system may be able to infer that a known type of attack is underway. Once such a state is established, predefined protective measures or damage-limiting measures could be invoked. Also, concomitantly, a knowledge base of known types of attacks and their signatures could evolve and serve as a knowledge source for the designers of survivable systems.

Ephemeral Entry Nodes

A distributed network–based system often has a designated entry node that may also be mirrored to ensure scalability. Malicious attacks on the system also enter through the same nodes. Changing entry nodes randomly or according to some secret algorithms could frustrate the efforts of the attackers who will try to enter through degraded or disabled nodes.

Conclusion

In this paper, we have outlined a number of design imperatives for the designers of survivable networked systems to consider. It should be noted that not all of these imperatives are applicable to every type of system. Some may be even in conflict with others and therefore may not exist simultaneously in the same system. Undoubtedly, these design strategies will significantly increase the complexity of the system and may degrade the performance of the deployed system. As with all systems, a careful trade-off between enhanced survivability and performance degradation should be considered. Emerging technologies such as storage-area networks (SANs) are addressing many of the issues outlined in this paper to make the future networked systems more survivable.

Template-Based Network Management, Ordering, and Service Provisioning

John Reilly
Principle Member of the Technical Staff
MetaSolv Software

Introduction

The rapid evolution of network technologies enables service providers to deploy networks that offer a variety of new services. Many service providers differentiate their services by constructing what resemble customized networks. This places an adaptability requirement on applications that support network management, ordering, and service-fulfillment business processes. Service providers must be able to bring these new services on-line quickly. This speed helps to ensure the provider's competitive advantage and allows the provider to quickly realize additional revenue associated with the new services.

This article describes the type of next-generation applications that satisfy this requirement. The applications should employ some type of *user-defined* templates that allow logical networks to be defined and designed. The templates should also contain rules that facilitate the ordering and provisioning of services enabled by next-generation networks. Functionality must also exist to allow users to define attributes unique to the networks and services offered.

What Is a Template?

A template is a collection of building blocks and rules that govern the design and provisioning of service-provider logical networks and the design and provisioning of customer networks and services (connections) that utilize the resources offered by provider's networks.

Four types of building blocks make up a template:

1. Types of networks
2. Types of elements
3. Types of connections
4. Custom attributes for each type

Types of networks specify the kinds of networks that make up a service provider's infrastructure and the kinds of networks offered to customers. Network types are the foundation building block for a template in that they are created first. Examples of service provider's networks include digital loop carrier (DLC), Ethernet, asynchronous transfer mode (ATM), digital subscriber line (DSL), multiprotocol label switching (MPLS), and so forth, as defined by a service provider's network planners. Examples of customer networks include virtual local-area network (VLAN), transparent local-area network (TLAN), virtual private network (VPN) intranet and extranet, and so forth, as defined by a service provider's network planners and product managers.

Element types define the groupings of equipment types that enable connections to be made within the provider's network or a customer's network. Element types represent the connection points in a network design template and are created next. For example, ATM switches, DSL access multiplexers (DSLAMs), frame relay (FR) switches, subscriber management systems (SMSs), routers, and switch routers are all examples of the types of elements that make up a network. Many types of networks can share the same type of element. For example, a switch router can be used as an element type in an Ethernet access network template, a VLAN, and an Internet protocol (IP) VPN. Element types can also represent other types of networks. For example, an ATM access type of network can be part of a DSL type of network.

Types of connections represent the manner in which elements can be linked. Connection types can be service-provider– or customer-related. Connection types represent the glue that holds a network together. For example, service-provider types of connections include bandwidth circuits, inverse multiplexing groups, facility circuits, Ethernet links, and so forth. Customer types of connections include bandwidth links, switched virtual connections (SVCs), permanent virtual connections (PVCs), Internet connections, TLAN connections, and so forth.

The network, elements, and connections that make up a network are described by very few hard facts, other than facts such as network name and network-element ID. Custom attributes represent additional facts that can be gathered and kept about a network, element, or connection. This is a necessity in today's rapidly changing technological environment. The attributes are defined and related to the template building blocks. When networks, elements, and connections are defined, these attributes are assigned values. For example, a TLAN ID is an attribute that represents the identity of a TLAN type of network. When customer X's southwest TLAN is designed, the TLAN ID is assigned a value of SW1.

Defining Template Rules

Template rules define how logical networks can be constructed, how orders can be placed for customer networks and connections, how customer connections are provisioned, and definitions of custom attributes and associations between custom attributes.

Logical Network Design Rules

Logical network design rules govern the construction of the networks. As seen in *Figures 1* and *2*, some of the rules can be graphically depicted. These rules include the types of elements that can be used to construct the network, along with the types of elements that can be connected. Definition of the rules does not end there. Rules also include, but are not limited to, how many of a specific element type can be used in a network, what type of equipment can be associated to each element type, what roles (such as customer edge, provider edge, and so forth) an element can play in a network and what type of connections are used to link elements.

Ordering Rules

Ordering rules facilitate ordering both customer networks and customer connections. The rules include linking items in a provider's product catalog to types of networks, types of elements, and types of connections in the network design templates. The rules also include such things as necessary element types that must be ordered to construct a customer network, location requirements for elements and connections, and time-of-day and recurrence sensitivity for connections. As an example of the latter rule, a multicast connection may be ordered for the next five Fridays between 2:00 pm and 3:00 pm.

Provisioning Rules

Provisioning rules govern the way customer connections are designed and activated. Some of the provisioning rules defined in a logical network design template are depicted in *Figure 3*. The rules in the figure are for a customer LAN to customer LAN virtual connection. The circled element types represent the types of elements used to design the connection. The numbered lines represent a possible path that the connection can take through the network of this type.

Some of the other rules not depicted in the figure, but included in the template, are connection assignment/allocation rules, equipment assignments, and whether the connection is a candidate for auto (sometimes called frictionless or touch-less) provisioning. An example of a connection allocation rule is the one stating that an Internet connection type must be assigned to the next available virtual path identifier associated to the bandwidth type of connection to which the Internet connection can be allocated.

Custom Attribute Rules

Custom attributes provide a unique way in which to extend the definition of network, element, and connection types. The rules define the custom attributes and relationships among custom attributes. Definition rules include the valid values for an attribute, applicable languages, whether the attribute is mandatory or optional, its data type and length, associated processes where it appears on a user interface, and its display properties. Relationship rules include the

FIGURE 1

A Logical Network Design Template Graphically Depicts What a Network Template Could Look Like

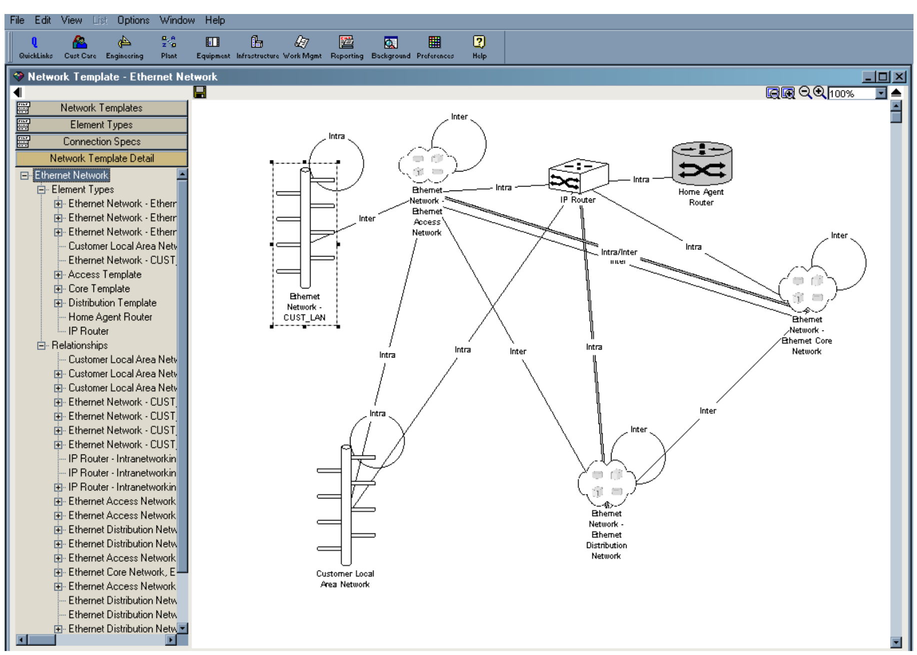

FIGURE 2

Logical Network Design Template Rules

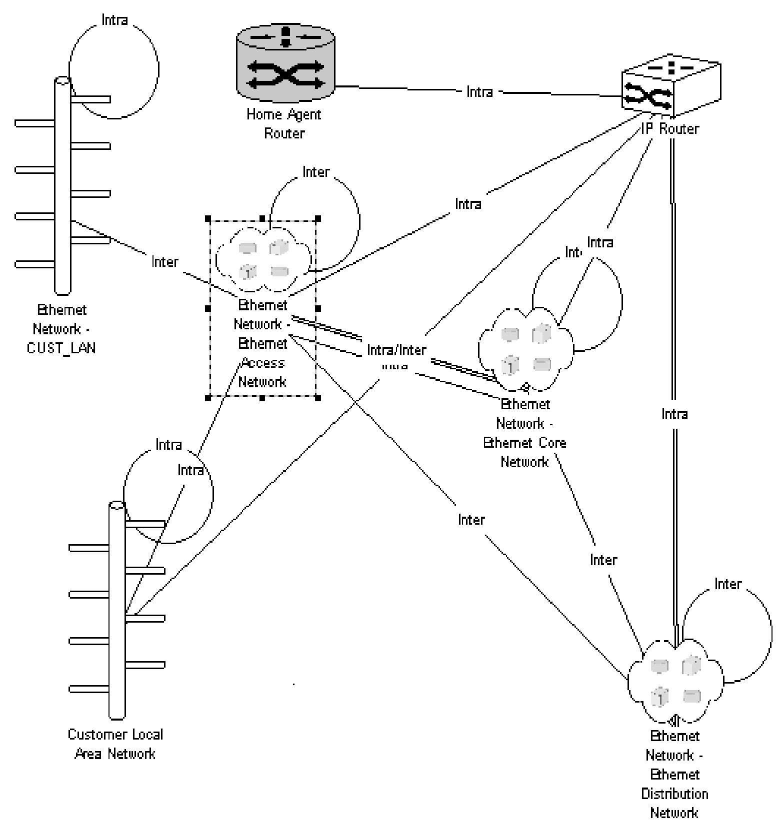

FIGURE 3

Template Provisioning Rules

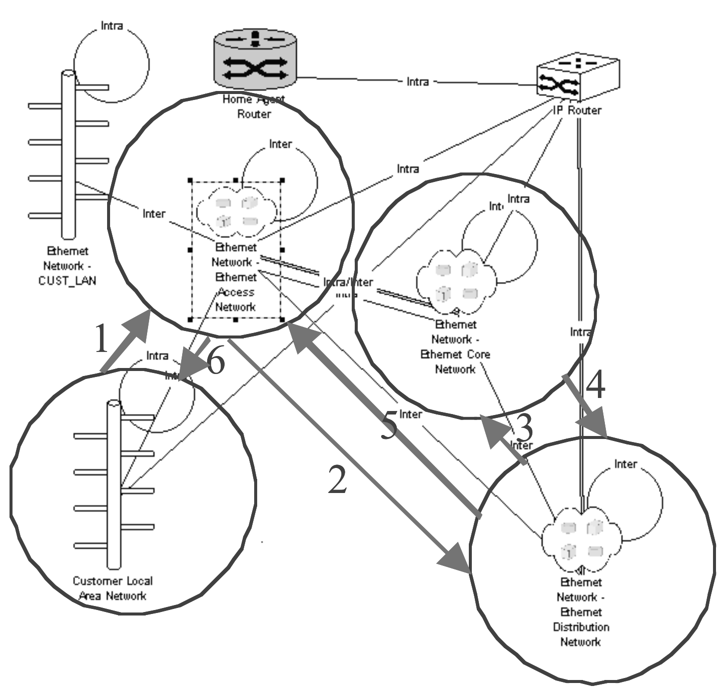

ability to default one attribute from another, validation that must be performed between attributes, derivation of an attribute from other attributes, and so forth. There should also be a provision for an attribute's value to be determined by some externally defined algorithm.

An example of the definition of a custom attribute is that its name is VLAN ID, that it is a number that can be a maximum of nine digits with no decimals, and that it is mandatory for a VLAN type of network.

One example of relationships between custom attributes is that the downstream bit rate defaults from the upstream bit rate. Another is that the committed information rate (CIR) is compared against bit rate for capacity tracking purposes.

Using a Template

Once defined, templates support a number of processes, including the following:

- Designing an internal network
- Accepting an order for a customer network
- Designing an ordered network
- Provisioning an ordered connection

Designing an Internal Network
Figure 4 shows a graphical representation of an internal (service-provider) logical network. This network's design is based on a logical network design template. Networks, their elements, and their connections are instantiations (occurrences) of network design template building blocks. For example, the northeast ATM access and southwest ATM access network are instances of an ATM access network template's building blocks.

The creation of the logical network was accomplished by completing the following steps:

- Selecting the template that will be used to instantiate the network
- Selecting the types of elements and instantiating them as elements
- Linking elements connected in the network
- Specifying the types of connections that link two elements
- Designing the connections that link the elements (optional if already designed)
- Associating connections with the link between two elements

Accepting an Order for a Customer Network
Relating items in a service provider's product catalog enables a process called assisted ordering. An assisted ordering scenario is contained in *Figures 5* and *6*. Once the type of network is known via its association to an ordered item, all of the ordering rules in the product catalog and logical network design template can be brought to bear.

In the first step shown in *Figure 5*, rules specify that site-type elements make up a VPN intranet. Assisted ordering, knowing this, can ask, "how many sites?" Locations are required for this type of element, as specified in the template. Therefore, locations are requested. Equipment is also associated to the element types in the catalog, so assisted ordering can ask about equipment to order for each site.

FIGURE 4

Internal Logical Network

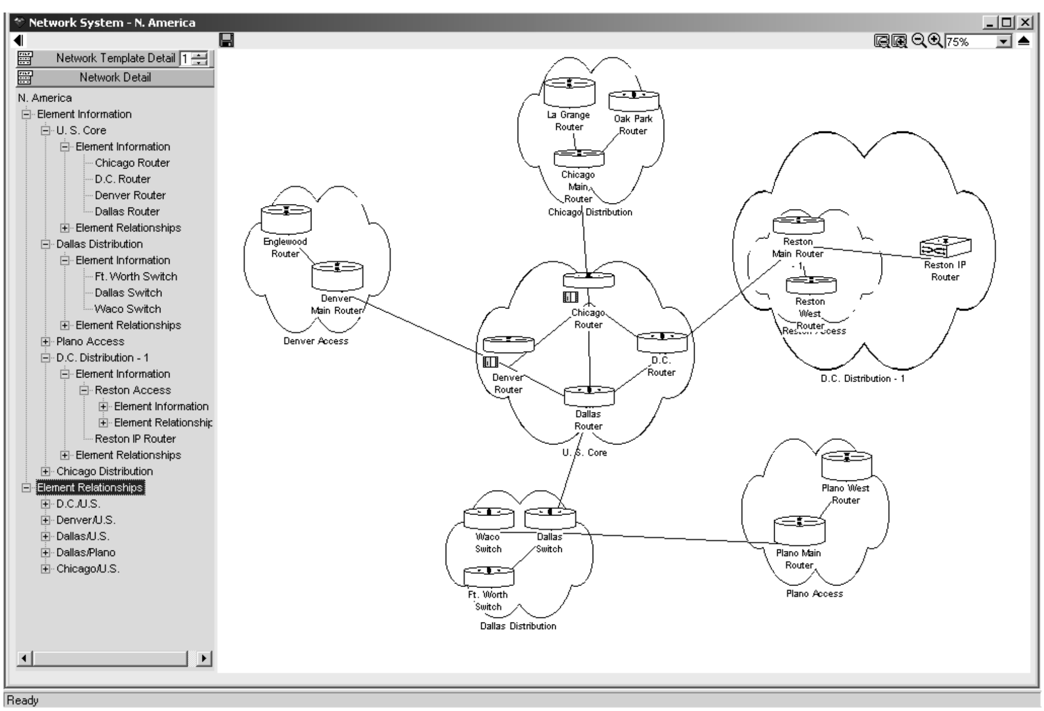

FIGURE 5

First Step of Assisted Ordering

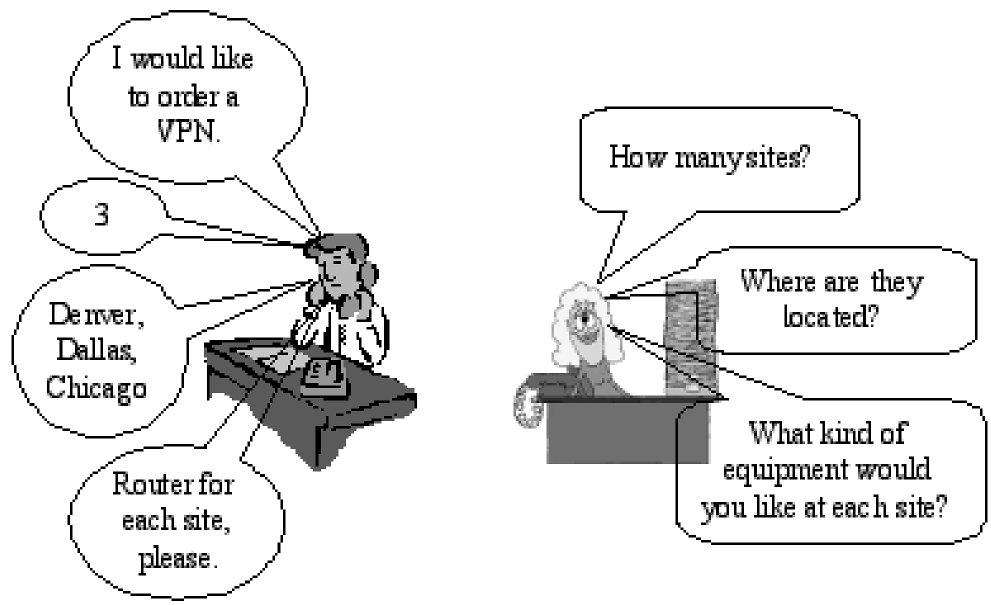

Similarly, ordering of connections can be assisted, based on the association of connection types to product catalog items that are part of the VPN intranet type of network. Using connection-provisioning rules, assisted ordering requests information about not only the types of connections desired, but also information about the terminating points of the connections. Use of these rules is demonstrated in *Figure 6*.

Designing a Customer Network

After accepting an order for a customer network, the customer network is designed as shown in *Figures 7* and *8*.

The rendering process uses information from the order and template rules to make a first pass at designing the network, as shown in *Figure 7*. Elements are placed on the drawing along with the serving portion of the service provider's network that will support the customer network, if the serving portion of the network can be determined.

The next steps of the customer network design process include the following:

- Linking elements connected in the network
- Designing the connections that link the elements (optional if already designed)
- Associating connections with the link between two elements (should be done automatically based on rules in the template)

FIGURE 6

Second Step of Assisted Ordering

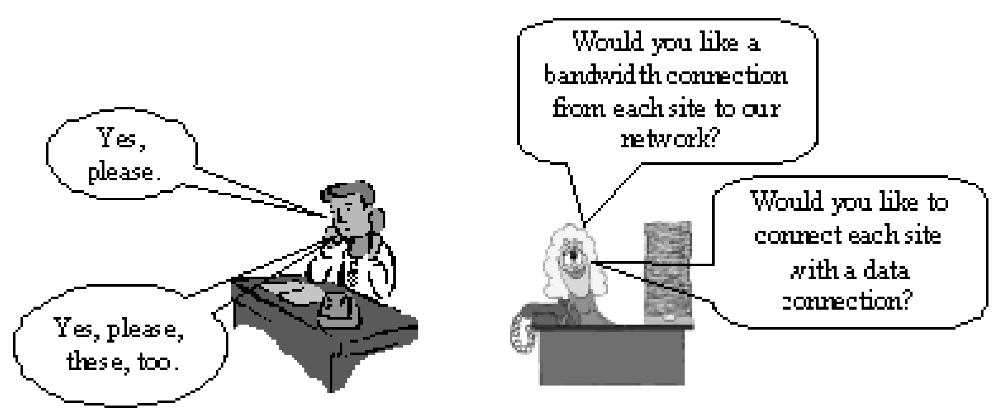

FIGURE 7

Initial Rendering of Logical Customer Network

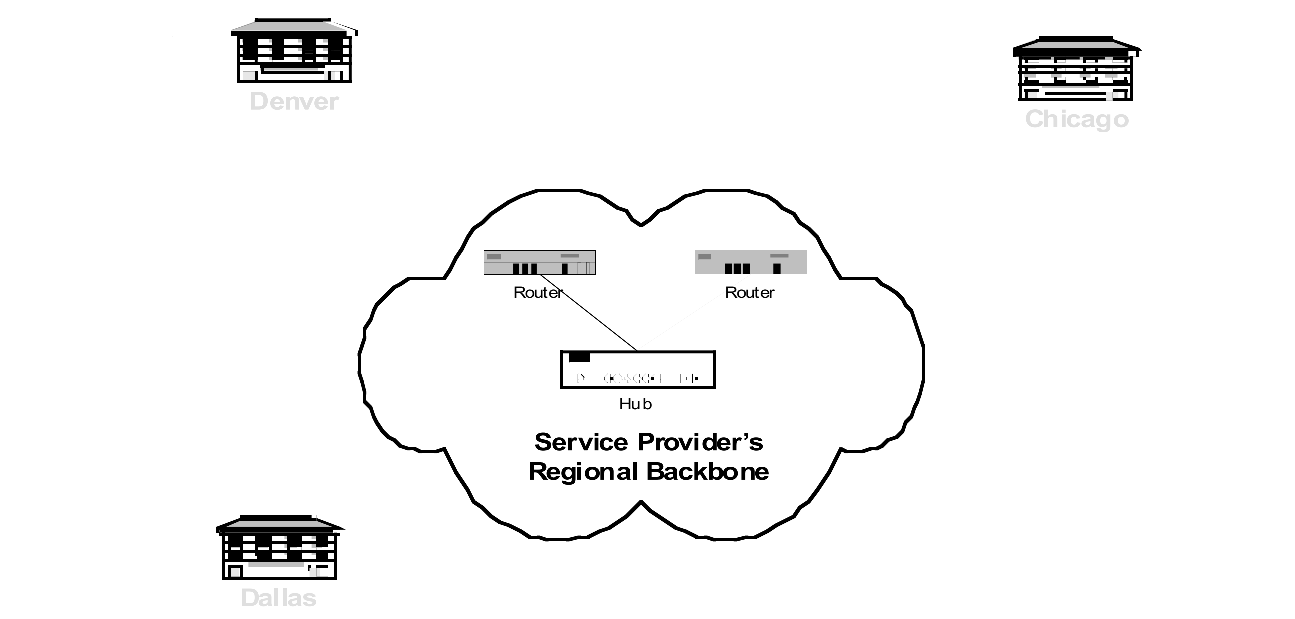

Provisioning an Ordered Connection

The design and activation of customer-ordered connections that are part of a customer network or standalone customer connections also benefit from rules contained in network design templates.

Figure 9 depicts a path through a service provider's network taken by a customer-ordered standalone connection between the two igloos. Finding a path is facilitated by rules in the template as described in the section on *Provisioning Rules*.

Design and activation processes are manually performed, totally automated, or partially automated based on rules in the templates and user choices made during the manual design process. For example, there are rules that govern the connection assignments and equipment assignments for each segment in the igloo-to-igloo connection shown in *Figure 9*.

FIGURE 8

Final Rendering of Customer Network

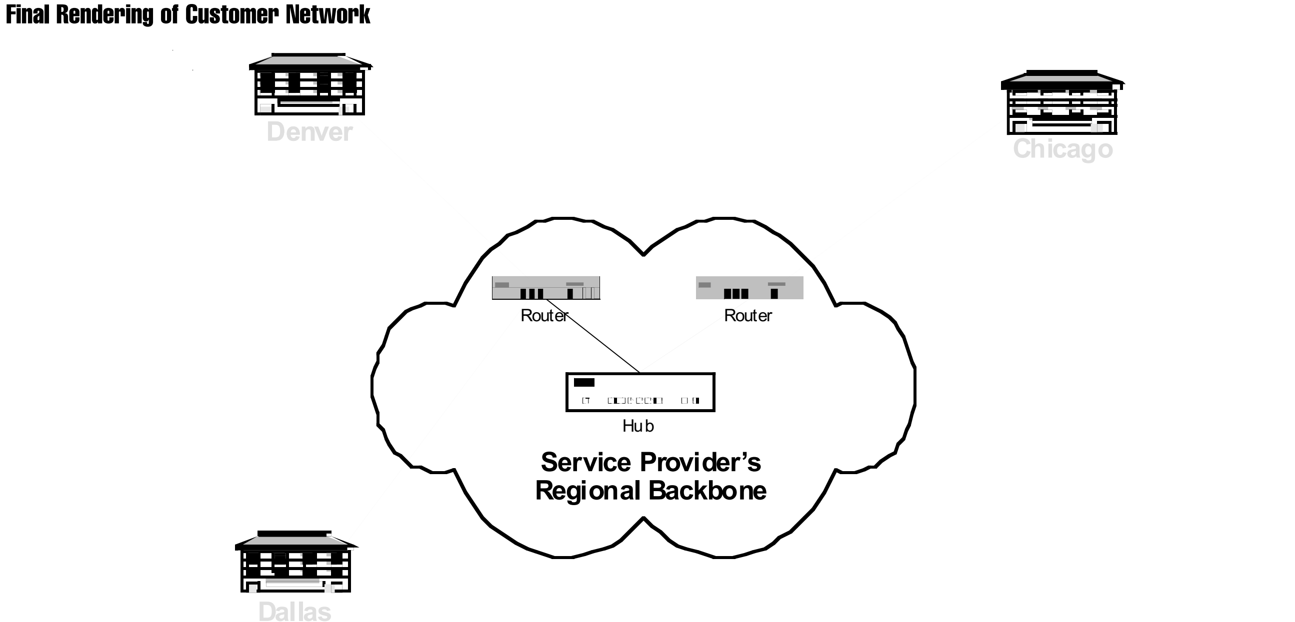

The steps involved include the following:

- Finding a serving network over which the connection will be provisioned
- Finding a path through the serving network
- Making assignments/allocations to other connections
- Making equipment assignments
- Activating the connection

Summary

Template-based network management is the wave of the future. Employ applications that are template-based to keep your competitive advantage.

FIGURE 9

Provisioning an Ordered Connection

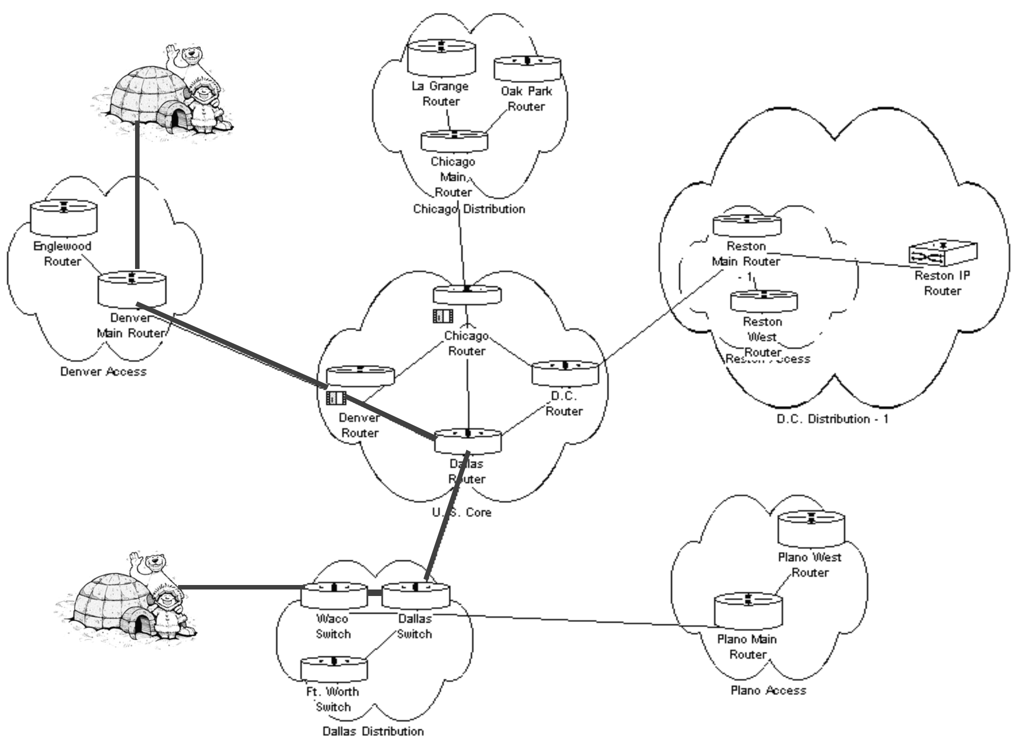

Automate Your Revenue-Assurance Organization

Gary Ross
Vice President, Product Management and Marketing
Connexn Technologies

In an age of automation, revenue assurance has strangely been left behind. It is ironic that while automation is a primary weapon in the war against revenue leakage, comprehensive automation of revenue recovery itself appears to have been passed over. It is true that some targeted areas of revenue-assurance automation have been addressed—for instance, call-through testing and usage-data monitoring in mediation operations support systems (OSSs)—but for the most part communications service providers' (CSPs') revenue-assurance organizations have few automated tools with which to work. With revenue leaking at the rate of 11 percent of annual gross revenues, this is a critical shortcoming that may very well affect a CSP's ability to survive in our current unforgiving economy. Or looking at it another way, state-of-the-art revenue-assurance tools may be the competitive edge that allows a CSP to pull ahead of its competitors.

In this article, I will describe the current dilemma for revenue-assurance organizations and suggest a way for them to overcome it. This is a critical issue that should be of interest to executive management working to improve margins by increasing top-line revenues and decreasing operating costs. I will make it clear why automation is critical to the process of revenue assurance itself.

Evolution of Revenue Assurance

First, let's take a moment to review the typical evolution of a CSP revenue-assurance organization to better understand how "lack of automation" happened and what needs to be done about it. Unfortunately, attention to revenue assurance usually begins with a catastrophic event, such as a switch that stops recording usage data, often resulting in multi-millions of dollars in lost revenue. Such a loss gets the attention of the executive team, which is quick to feel the pain and prone to action. Resources are dedicated to clean up the mess and ensure that it doesn't happen again. An executive governance or oversight committee may be formed, and under their direction a revenue-assurance organization is created, frequently in the finance organization.

This newly formed and empowered revenue-assurance team sets out to identify risk and develop initiatives to protect against leakage. At this point, they begin to build staff and alliances, either centralized or distributed, throughout operations organizations. Overwhelmed with the enormity of their task, for revenue-assurance is an end-to-end concern, they soon seek out ways to be effective, and this is where their troubles begin. Their options are limited. Even though empowered by the executive team, the only way to keep up appears to be by adding staff. Their work is labor intensive and frequently involves time-consuming manual audits. Consultants are more than happy to roll up their sleeves and pitch in, but their contribution is mostly manual and at high cost. Internal information technology (IT) organizations are sometimes willing to take on point solution automation projects, yet frequently are a dollar short and a day late. This is where revenue assurance gets bogged down. The initiatives stack up, and the organization gets stressed.

Clearly, manual revenue assurance and point solution automation projects simply are not an option in light of the ever-increasing complexity of CSP revenue streams and the sheer volume of work to be done.

Revenue Leakage

At this point, it will be helpful to examine revenue leakage itself, which will give insight into why automated tools are a must. There are really two kinds of leakage occurring—one that I refer to as *systemic* and another that I refer to as *pervasive*.

Systemic leakage impacts many customers of which a single occurrence can result in huge sums of loss. An example would be a network-element translation error that turns off usage recording for many customers. This is also the kind of leakage that is most visible, referred to as low-hanging fruit by some. To date, most point solution attempts at automation have been focused on systemic leakage, such as "call-through testing." In addition, revenue-assurance metrics and controls are also focused on systemic leakage. In other words, there are some clear-cut ways to go after systemic leakage, and one might argue that mature revenue-assurance organizations are doing a pretty good job of it.

FIGURE 1

Service Providers Today

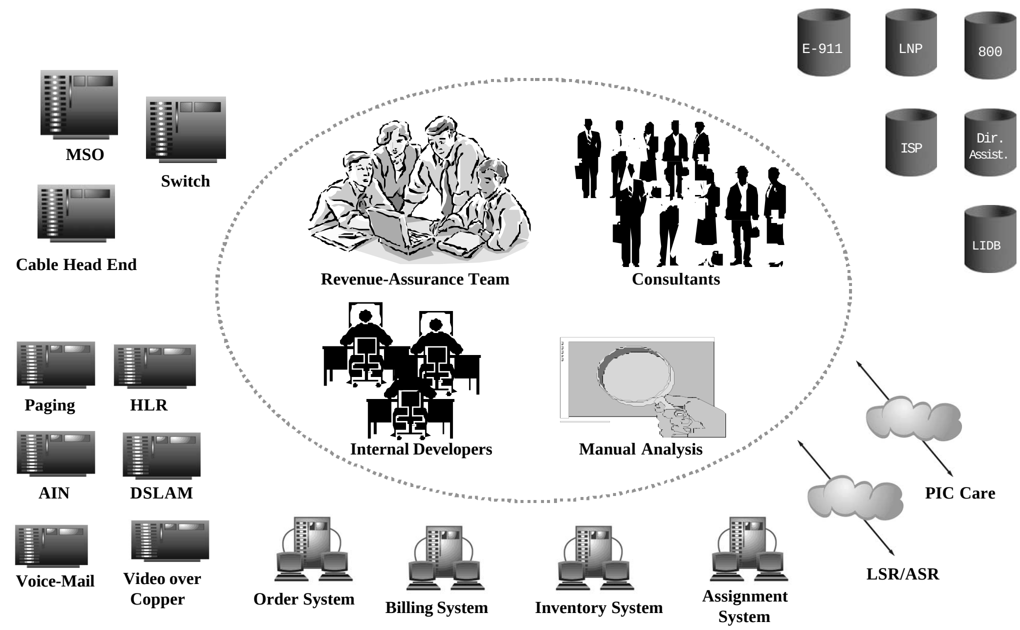

Pervasive leakage, on the other hand, is associated with a single customer or network component assigned to a single customer. An example is the customer who has been provisioned for voice-mail but is not being billed for it. Pervasive leakage takes the form of unbilled/underbilled services and features or assets that have been lost track of (stranded assets) and therefore are no longer earning revenue. Pervasive leakage is really tough to go after, since it usually requires accessing and analyzing huge amounts of data. It also usually requires the comparison of data from two or more sources. Not only is the data likely to be difficult to access, but it is also rarely, if ever, in a comparable format. Anyone who has done services and features audits, manually or otherwise, has found astonishing rates of underbilling of lines across their customer base. Connex^{n} Technologies has found that 3 percent to 20 percent of all lines/subscribers are being underbilled, with an average of 12 percent. Most CSPs suspect some level of pervasive leakage, but they usually have no hard data on how much. Connex^{n} Technologies has also found that most CSPs are not even looking for pervasive leakage, mostly because of the time and resources required to detect and correct the errors. Yet, ironically, this is where the highest payoff of revenue-assurance work is, and today it simply is not being done. Enter the need of automation.

Today

Today, most mature revenue-assurance organizations are slowly working through long lists of initiatives at the slow rate allowed by resources and funding levels. Metrics and controls have been put into place. The low-hanging fruit has been harvested. The tough work of addressing pervasive leakage has now begun and the going is extremely slow. Their instincts tell them that unseen revenue must be leaking, yet they have no hard data to substantiate their feelings. Worse yet, they have no capacity to go after it. Key to their long-term success is access to data locked away in business processes and OSSs. Looking for pervasive leakage requires access to volumes of data and the ability to efficiently and effectively audit for leakage. *Figure 1* depicts the typical revenue-assurance organization and the lack of automation that they face today. The data of value for revenue-assurance work is distributed widely across most corporate systems and processes. Access to the data usually requires a custom interface. Once the data is accessed, it then needs to be transformed before it can be compared to another data source.

Tomorrow

What revenue-assurance organizations need is a flexible, automated way to access and analyze data from distributed sources. In addition, they need an automated approach to the correction of their findings. In short, they need a revenue-assurance support system. *Figure 2* depicts a revenue-assurance organization with a revenue-assurance support system. In general terms, such a system will need three levels of functionality, all three of equal importance. First, it must be able to connect to a wide variety of systems and network elements. The connections must be reasonably easy to establish and maintain. Second, it must have the flexibility to interact with those systems in a way that allows it to both extract desired data and to communicate changes in the data back to those systems. This will facilitate ad hoc audits and allow for automated error correction. Finally, a revenue-assurance support system must have robust reporting capabilities that allow quantification of audit results both in terms of revenue and root cause.

FIGURE 2

What Service Providers Need

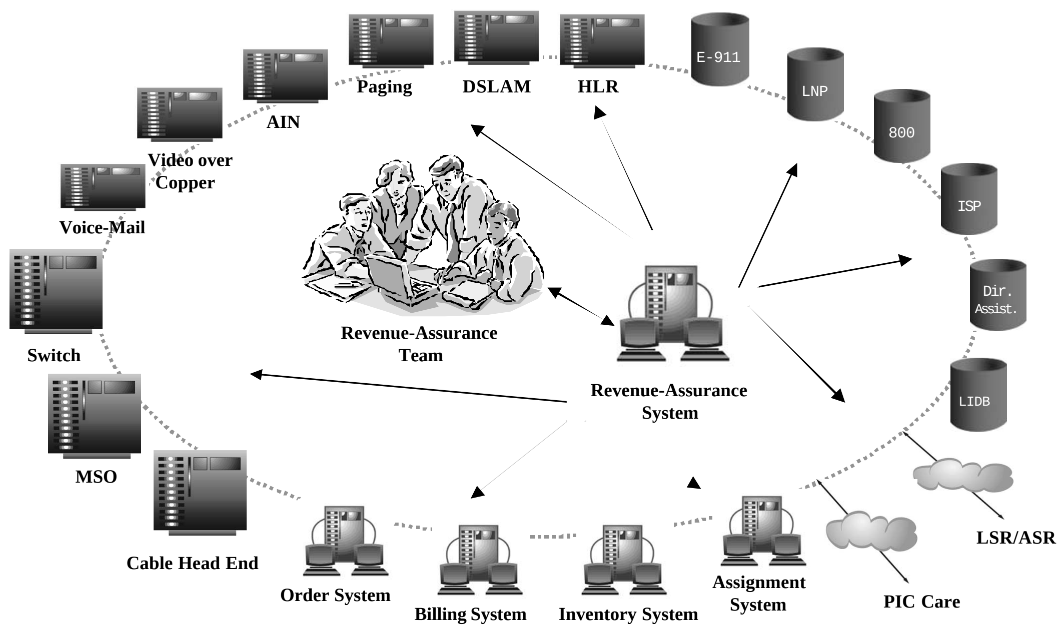

The obvious question at this point is does such a system exist? I would say yes, in its early stages. Automated detection of pervasive leakage and automated correction of the findings has been going on for some time. The ability to connect and communicate with a wide range of systems and databases is also available. However, the idea of a user interface that supports ad hoc auditing and performs advanced reporting is somewhat new and will be evolving over time as CSPs clarify needs and requirements.

Conclusion

Automation is necessary for revenue-assurance organizations to bust through current resource constraints. During a period of economic downturn, when capturing all revenues earned is more important than ever, manual auditing and analysis is simply not an option. CSPs need to realize that this is why their revenue-assurance organizations are stalling out and take action. Now is the time to take revenue assurance to the next level.

Roots of Reliability

John W. Sawyer
Director, Mission-Critical Facility Services
Johnson Controls, Inc.

On-line business is driving explosive growth in computer networks. These high-value systems are destined to fail without proper attention to the basics of climate control and power quality.

"(A) popular online auctioneer...is still working on a major system problem that has disrupted service since last night. (The) latest outage is its second major service disruption this week alone, and continues a string of service outages that have plagued the company since the beginning of this year.…Though the company didn't offer details, earlier bulletins blamed a hardware issue relating to (the company's) database servers."

—*Computerworld*, June 11, 1999

No matter where, when, or how long, network outages make news. They are the nightmare of on-line merchants, Internet service providers (ISPs), telephone and wireless communications companies, banks, and on-line brokerages.

An outage means immediate loss of sales and revenue. It also can mean loss of customers, a steep decline in market capitalization, and lasting damage to the company's reputation.

The cause of a given outage can be as obvious as a lightning strike or as arcane as flawed application software code. But often, the root cause can be traced a flaw in basic infrastructure: the electric power and climate controls that support critical computer equipment.

What shuts down an on-line merchant at the height of holiday shopping may be a remarkably low-tech problem:

- A power failure traced to an uninterruptible power supply with insufficient redundancy or imbalanced load
- A server room that ran too warm because an improper rack configuration blocked the flow of conditioned air
- An inadequately trained employee who tripped the wrong circuit breaker during a routine maintenance inspection

Recent events on Wall Street show the financial perils of new-economy companies growing too fast, too soon. Such companies can place themselves at similar risk by building critical infrastructure with an eye on speed to market but without systems in place to ensure uninterrupted service to customers for the long term. Just as lasting success on the stock market depends on sound business fundamentals, reliable critical equipment depends on sound engineering and operating standards at each step of development:

- Site selection
- Design
- Construction
- Commissioning
- Steady-state operation

"I think you need to bring an elephant gun to kill a mouse. Hardware is cheap, the pace of the game is frenetic, and being the first mover in an industry like this is very important....You really need to figure out what the business plan is, do a what-if scenario that is beyond your wildest dreams, and build an architecture that lets you scale beyond your wildest estimation."

—Maynard Webb, Chief Information Officer, eBay Inc.
Computerworld, September 8, 1999

Critical Facilities: Brave New World

Critical facilities in business are far from new. Banks, telephone companies, credit-card merchants, and brokerages have owned vast computer networks for years. What has changed with the new economy is the speed of deploying technology.

Since the Internet emerged as a communication channel and a place to shop for goods and services, demand for critical facilities has exploded. It is estimated that within the next 36 months, annual sales through e-commerce will reach $60 billion in the consumer sector and five times that for business-to-business (B2B) (see *Table 1*). Consumer transactions will increase 22 percent, and B2B transactions will increase 78 percent.

This growth profoundly affects the telecommunications, Internet service, and fiber-optics industries. New and existing companies, jockeying for market position and points of presence (POP), rush to build new critical facilities, with speed and budget as their key criteria.

Those are far different from the factors that drive critical facilities development in more established industries. Ten years ago, the typical critical facility was a data center that housed mainframe computers and a range of peripheral devices. Today, those companies still prefer to concentrate their network and data technology in one or a few locations.

TABLE 1

Projected Growth in Electronic Commerce: 2001 through 2003	
Sector	**Increase**
Financial	400%
Personal Computers	340%
Travel	1,031%
Entertainment	806%
Apparel	460%
Books/Music	605%
Tickets/Events	2,432%
B2B	2,250%

To insulate against risk, they grow according to a plan, and they apply strict engineering, operation, and maintenance standards. Support personnel are experienced and technically capable. In-house staff, architects, and consulting engineers form a network of expertise, constantly vigilant for trouble. They are familiar with clients' business processes and the associated business plans and projections.

That stability and discipline are often missing in the competitive world of the new economy. Fledgling technology companies financed by initial public offerings (IPOs) staff up quickly, hiring people from multiple organizations. Critical facilities staffs are hybrids of people with facilities, information technology (IT), and management backgrounds.

The speed-to-market imperative greatly accelerates the build cycle. At the same time, cost pressures dictate the maximum productive use of floor space. Equipment configurations may violate temperature and humidity requirements and may fail to consider future maintenance needs. Facility design and operating practices may lack standardization across locations.

Most serious, growth may outstrip the capability of standby electrical power, power quality protection, and air-conditioning systems. While the occasional power or climate-control problem may not seriously harm a traditional business, it can devastate an Internet-based firm, especially one that has lured customers with guarantees of nearly 100 percent reliability.

The risks reach well beyond Internet companies themselves. Today, many corporations outsource specific computer software applications and related support by using Internet-based application service providers (ASPs). However, transferring technology means transferring risks of network failures and downtime, along with related consequences.

"More than 1,000 companies this week experienced slow performance and other problems on their Web sites because of troubles that (their) Internet hosting firm had with a router at its data center...."
—*Computerworld*, June 28, 2000

"If you find out your system can't handle the job, it's not something you can buy your way out of on the spot. You have to plan ahead."
—Stephen Scullen, Fidelity Brokerage Services Inc.
Computerworld, April 10, 2000

The Anatomy of Risk

Critical facilities must closely regulate temperature and humidity and their rate of change. Operating outside manufacturers' specifications can decrease mean time between failures (MTBF) in servers and other computer data storage and communications devices. Low humidity increases the risk of damaging electrostatic discharge.

In recent years, manufacturers have made computer equipment somewhat more temperature and moisture tolerant, but 68°F to 72°F and 40 to 60 percent relative humidity are still considered ideal conditions in critical facilities.

Critical equipment rooms are typically cooled through subfloor plenums that deliver conditioned air under pressure through a system of floor vents. In a well-designed system, cooling capacity is sized to the equipment's heat-rejection requirements, and air pressure and flow patterns are carefully managed to deliver the necessary cooling to each area of the room. For example, floor openings are strategically located to discharge air in front of servers so that the units' cooling fans can draw it in.

The overall room design recognizes manufacturers' requirements for "white space"—open areas around racks that enable front, back, or side access for replacement and service.

Problems occur when companies strive to maximize revenue potential per square foot of space outside the context of a sound facilities plan. When servers and other devices are added without considering the effects on air conditioning and power requirements, the critical environment can slip out of control.

Often, personnel building critical equipment rooms ignore "white space" requirements, installing equipment in any open area of floor. Besides increasing heat load, the added equipment has a side effect, requiring more electrical cable cutouts in the floor. These allow the uncontrolled escape of conditioned air, reducing subfloor pressure and impeding the delivery of air to points where it is most needed. Further, improper alignment of equipment racks can impede return airflow, causing air stratification or "hot spots."

As cooling effectiveness declines, a common reaction is to install more air-handling units in the room. At best, this is self-defeating, because those units consume the same valuable floor space from which the company seeks to maximize revenue. At the worst, in extreme cases, equipment-cooling requirements exceed the capacities of chilled water piping, system air volume, condenser loops, or cooling towers. This easily happens because system efficiencies, age, and maintenance histories are seldom taken into account.

In any case, the net result is loss of the ability to cool equipment properly over an extended period. Equipment runs above manufacturer-recommended temperatures, or, just as serious, its temperature fluctuates beyond manufacturers' rate-of-change specifications. While this may not immediately affect reliability, it inevitably accelerates wear and in the long run increases the risk of failures. High outdoor (ambient) temperatures aggravate these conditions.

Ill-planned commissioning or installation processes also can affect electric power reliability. For example, improperly organized circuits may create phase imbalances large enough to cause an uninterruptible power supply (UPS) to trip off-line. At bare minimum, a level of power-protection redundancy is lost, and the risk of a system outage is magnified.

"In response to severe web service outages, (a web hosting firm) plans to offer affected customers one month of free service, worth about $150."

—*Computerworld*, July 10, 2000

"Giving credit for downtime is not what these people want. They want reliable service."

—Rosemary Cochran, Vertical Systems Group
Dedham, Massachusetts

"The lesson there, I guess, is that I would gladly pay a hundred times what I got refunded not to have the outage."

—Charles Rice, Coatingsmart.com Inc.
Ann Arbor, Michigan

Cases in Point

A properly designed system, an ongoing management program, and properly trained personnel are keys to effective critical environment control. Two recent examples illustrate how the risk of failure increases when either component is lacking.

Case 1: Flawed Design
A major computer software and on-line service company was experiencing climate-control problems in a 15,000-square-foot critical facility. As the company installed critical equipment in the raised-floor room, temperatures were slipping outside specifications in multiple areas. In response, the company commissioned an analysis of the climate-control system, including air handlers, the condenser loop, and subfloor airflow. The analysis found the following:

- The 15-year-old condenser loop was providing insufficient heat transfer because of severe tube scaling caused by inadequate water treatment.
- Air-conditioning equipment was running at full capacity with efficiencies of less than 60 percent and with multiple equipment problems—yet no alarms were being produced.
- Improper alignment of server racks impeded the return path of air to the in-room air-handling units, causing hot spots and stratified air conditions.
- Obstructions in the subfloor lowered the available subfloor air pressure to near zero in certain areas, preventing delivery of conditioned air.

In summary, the facility was at significant risk of a failure that might have caused a system outage affecting large numbers of customers and possibly leading to widespread adverse publicity.

The company corrected the problem by upgrading cooling equipment, moving the air-handling units so that return air flowed perpendicular to the units, removing subfloor obstructions, introducing a process to manage floor cutouts, and installing chemical treatment programs. The modifications made temperatures in the room far more uniform.

Case 2: Training Deficiencies
Studies of critical facility service interruptions document that eight of 10 are caused not by the initial event (such as a lightning strike or dig-in fault), but rather by how systems or personnel respond.

An Internet-based services organization faced a potential crisis in its call center after one of two UPSs in a redundant system tripped off-line. Immediately after the event, operating personnel called for service. The service technician assessing the problem went to the circuit breaker to disconnect the UPS that had failed.

Without consulting the as-built drawings, the technician threw the breaker labeled for the failed UPS. In reality, the labeling was incorrect because it was not changed after a recent retrofit. As a result, the technician's action cut the utility power feed to the operable UPS.

That UPS went into alarm, reporting to the building automation system that its battery was discharging and that a shutdown of the protected systems was imminent. No one noticed the alarm. When battery power ran out, a major outage occurred, affecting all terminals in the call center.

The incident could have been avoided if labeling of electrical circuits had been correct, if the technician had not assumed the labeling was accurate, if facility personnel had been properly trained and drilled in handling such events, and if strict policies had been enforced on system monitoring during abnormal operating conditions.

"System outages and capacity problems rank as the chief concerns among executives at online brokerage firms, according to a new survey by Deloitte & Touche LLP. The New York–based accounting and consulting firm polled 60 online trading executives between November and January. It found that 60% of executives at full-service brokerages and 38% of managers at discount firms listed 'system outages and mistakes/handling growing transaction volume' among their fears."

—*Computerworld*, April 10, 2000

"We're reaching the stage where (customers of web hosting companies) are going to want to have telco-grade reliability. There's no excuse, really."

—Courtney Munroe, International Data Corporation
Quoted in *Computerworld*

Toward Long-Term Reliability

Reliable critical facilities depend on supporting infrastructure built with discipline and foresight. In the long run, a slower, planned approach speeds progress by eliminating errors and resulting rework. In the bargain, the risk of outages during system evolution declines dramatically. An effective critical-facilities plan accounts for four key components: people, process, systems, and technology.

Building for reliability is a five-step process that allows no shortcuts. Each step requires coordination among multiple functions. Those who design, engineer, and construct the facilities must consider the needs of those who will eventually commission, operate, service, and maintain them.

Site Selection: Construction cost per square foot is a key business consideration; low-cost real estate is thus attractive. But cost savings must not overshadow the long-term operating requirements of critical facilities. Needs like security, power quality, fire protection, and structural integrity may easily outweigh low cost.

Companies planning critical facilities should also consider local earthquake potential, lightning strike frequency, crime levels, and, most important, electric service reliability. Key considerations include the frequency of power outages (scheduled or unscheduled) and of switching on the grid. Sites near foundries or other operations with large electrical loads require careful planning for power-quality protection.

While older buildings can be appealing because of their low prices, their disadvantages can be considerable. Potential issues include roofing integrity, excessive numbers of windows (potential break-in points), floor load-bearing capacity, and absence of floor drains (essential for fire sprinklers). Some older buildings may be technically unsuitable; others may have renovation costs that erase their initial price advantage.

Design: Facilities must be designed around specific standards that cover equipment and material specifications, code compliance (fire, electrical, plumbing, ventilation, and others), and installation practices. These standards should be consistent from location to location and must relate to original availability requirements and staffing considerations. Critical considerations include the following:

- *Utility Power:* Historical grid reliability figures going back three years or more should be secured.
- *Cooling Capacity:* Cooling equipment—chillers, piping, duct capacity, headers, condenser loops, cooling towers—must be sized for the heat-rejection needs of the fully built facility, with allowance for system losses. Computer equipment manufacturers' cooling recommendations must be observed.
- *Airflow:* Equipment must be arranged to prevent airflow obstructions, which can create warm and cool spots and thus allow some equipment to overheat. Server racks should be aligned parallel to the direction of return airflow. Cable cutouts and other floor openings must be minimized to prevent misdirection of air and loss of air pressure. Subfloor airflow restrictions can be reduced with cable-management systems that direct electrical cables into troughs. Such problems can be nearly eliminated with overhead electrical cabling.
- *Maintainability and Safety:* Systems must be configured so that maintenance can be performed without sacrificing redundancy. For example, cooling system condensers should be connected in parallel (not in series) so that a repair to a single unit does not require shutdown of all or most units. Multiple power buses require adequate clearance so that one can be serviced without de-energizing the others.
- *Standardization:* Use of the same materials, devices, and technologies across all sites helps reduce costs in the design, construction, and operating phases.

Construction: Facility buildout must follow accepted installation best practices and must adhere to all design standards and metrics. When possible, systems and equipment should be tested to verify that they meet design specifications. For example, UPS systems should be factory witness tested under load. Tests should also be conducted for continuity in grounding systems, on links between generators and transfer switches, and on power-cable insulation.

Commissioning: Before start-up, design specifications and construction standards require verification by an outside source hired by and accountable to the facility owner. The entire power system must be confidence tested. Equipment warranties must be initiated. Commissioning also includes the generation of critical documents: as-built drawings, standard operating procedures, general maintenance plans, and training procedures. These, along with equipment specifications, warranties, and call numbers for service, should be placed in a reference library. Factory representatives who supplied the equipment should take part in staff operations and maintenance training. Baseline performance data and confidence-testing data should be captured. All critical procedures should be tested and verified and a training program established.

Steady State: Sound operation, maintenance, and contingency plans should guide day-to-day activity. Plans should include an escalation policy specifying whom to call (on-site staff, outside repair service, vendor representative, corporate technical support) for different levels of alerts or alarms. Critical spare parts should be identified, sourced, and stored on site as required. (In this regard, standardization on products adds economies of scale.)

Whenever a change in critical equipment is proposed, appropriate personnel should carefully evaluate the effects on the critical environmental. If the change is made, as-built drawings, maintenance procedures, training programs, and other facets of the operations plan must be updated. This process should include updates all associated document, equipment-labeling, and equipment-testing procedures.

A change management policy should guide any necessary adjustments to operating procedures. To protect "tribal knowledge" held by long-time staff, personnel should be cross-trained and succession plans should be in place for key staff members.

The Reliability Imperative

Without exception, technology companies are best served by designing, building, and maintaining critical facilities according to a comprehensive, integrated plan.

Companies seeking advice on critical systems design should look to an advisor with extensive experience in building critical environmental control. Such a company will be qualified to devise a cost-effective supporting infrastructure; to help assemble a coordinated team of company personnel to plan, build, and operate the system; and to provide financing options for critical assets.

In the long run, this approach allows the development of sound, reliable systems that meet customers' service expectations.

QoS Your Way

Sara Sedighi
Product Manager
Metro-Optix

Considering that most service providers currently have solid infrastructures providing the highest level of quality of service (QoS) with synchronous optical network (SONET) transport, how do they maintain this superior quality level and yet add new revenue sources as we migrate toward data-centric networks? With the explosion of the Internet and the high demand for enhanced data services, service providers are struggling with the challenge of this migration without stranding their current customer base and sacrificing the corner stone of their business. Most analysts project that data traffic will be the dominating network traffic in the near future, and service providers need to strengthen their position in the new data market in order to maintain their profitability without sacrificing reliability, survivability, and quality.

Transitioning from the legacy time division multiplexing (TDM) services to the next-generation data-oriented optical networks makes a compelling business case simply because data traffic is mission critical in today's network. But should the delivery of quality and reliable services be sacrificed along the way? Due to customers' demand for multiple service types, service providers are faced with the challenge of managing multiple network elements to support all of the different services and networks. Provisioning and maintaining all of these different network elements makes it prohibitively expensive for the service providers, potentially limiting their business opportunities due to cost constraints.

What is really needed is a single network element that provides support for multiple-tiered, multiple services and multiple networks without sacrificing quality and reliability. This ideal network element should also support QoS for each service layer. Fortunately, next-generation metro access vendors are deploying bandwidth managers capable of handling all these solutions—providing one-stop QoS shopping!

Bandwidth Managers

The quest for a complete QoS solution has introduced a new breed of next-generation bandwidth managers. Developers of these new bandwidth managers understand the importance of offering a variety of services with commitment to deliver QoS for each service layer. These new bandwidth managers are designed with built-in QoS support and the ability to support multiple protocols without requiring hardware forklift upgrades as customers' traffic patterns shift from traditional voice to converged voice and data. Support for different transmission methods (whether cells or packets) provides total flexibility and freedom in designing a reliable and efficient network.

And how do bandwidth managers achieve this level of flexibility and efficiency with optimal reliability and quality? They combine different protocols in one compact, high-capacity network element. They handle TDM, asynchronous transfer mode (ATM), and Internet protocol (IP) traffic with the highest quality to satisfy the needs of today's demanding and versatile networks—true protocol agility! These new bandwidth managers can easily scale from all–TDM to all–ATM to all–IP, or any incremental combinations of all three protocols (see *Figure 1*). This capability will help the network services provider to seamlessly migrate today's TDM network to the next-generation data-oriented optical network, while still providing traditional TDM transport. The flexibility of a bandwidth manager enables the service providers to provision any type of service based on their customer demand and their growth. Changing from TDM services to ATM and IP does not necessarily translate into inferior QoS. Here is how they address the QoS issue, their way!

TDM QoS

The SONET (TDM) portion of the bandwidth managers support the quality transport of Layer-1 traffic. The nature of a circuit-switched TDM network is to support dedicated connections and real-time data transfer, and that's what SONET networks offer—100 percent availability and reliability. The circuit-switched lines are still the best and most reliable for carrying voice traffic, therefore delivering the highest level of QoS. A bandwidth manager with SONET synchronous transport signal (STS) and virtual tributary (VT) multiplexing and grooming offers guaranteed bandwidth and a high level of service in the network.

ATM QoS

The ideal bandwidth manager provides QoS with no restriction on the service type, which enables the service providers to defer services and cost until the need arises. When service providers are ready to deploy ATM services in their network, the ATM portion of a bandwidth manager supports QoS at the ATM layer independently or simultaneously with other layers.

FIGURE 1

Bandwidth Manager

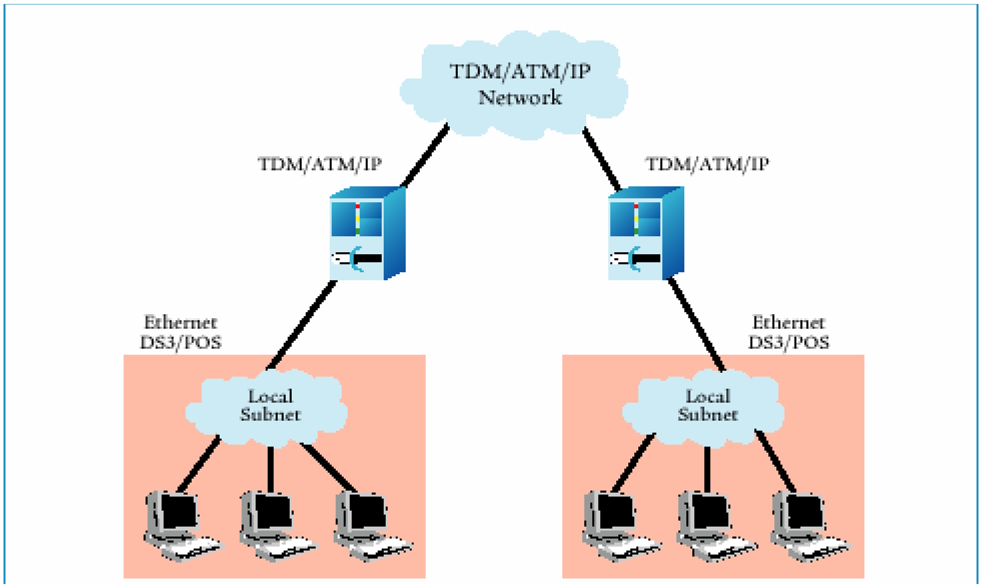

ATM is inherently connection-oriented and started as a single-queue traffic management, but it could not meet the multi-tiered, multiservices needs of today's networks. To provide multiple services with negotiable service-level agreements (SLAs) and maximize the performance of data networks, ATM QoS technology started to take shape and multiple classes of service were introduced. ATM QoS is implemented by defining five classes of service: constant bit rate (CBR), variable bit rate–real time (VBR–rt), variable bit rate–non–real-time (VBR–nrt), unspecified bit rate (UBR), and available bit rate (ABR). The service class defines parameters and acceptable cell loss and delay variation through the network. As each virtual connection (VC) in an ATM network is created, a specific service class is assigned to it. To provide complete QoS coverage, these bandwidth managers support all five traffic classes: CBR, VBR–rt, VBR–nrt, UBR, and ABR. Depending on which services are activated, their scheduling mechanism selects which VC queue to service at each cell dispatch time and whether to buffer ingress/egress data. This enables them to meet the QoS requirements of different VCs while simultaneously promoting high utilization of both bandwidth and buffers in the switch.

The services are divided into three levels of priority:

1. The highest priority is accorded to CBR and VBR–rt traffic. These services demand committed bandwidth and low latency. Bandwidth managers reserve enough bandwidth in each ATM node for CBR and VBR connections.
2. The second priority level is for guaranteed traffic services (GTS), which include the VBR–rt and VBR–nrt classes. These services are characterized by the average cell rate and are used for carrying packetized voice, video, and data. Within GTS, the different services are shaped according to what service customers agreed to pay for.
3. The lowest priority level is called excess bandwidth service (EBS) or best effort. This service is available in weighted shares for ABR and UBR as determined by programmable weights.

Service providers are able to offer QoS differentiated services based on the different service classes and varying bandwidth guarantees in the traffic contract. Service classes and traffic contracts are enforced using input policers and output schedulers that enable bandwidth managers to measure input traffic rates, identify the violating traffic, and manage the traffic with multiple queues and sophisticated algorithms to guarantee service classes.

The evolution of ATM QoS from best-effort services to guaranteed services has remarkably changed the QoS standards such that ATM networks offer QoS levels comparable to dedicated TDM networks, therefore providing maximum reliability and availability. This makes it possible for service providers to offer ATM QoS services independent of the underlying transport technology. Best-in-class bandwidth managers provide full support for all of these multi-tiered, multi-classes of service.

IP QoS

As service providers are preparing to meet the increased customers' demand for bandwidth with typical applications such as virtual private network (VPN), voice over IP (VoIP), and Web hosting, the need for a new set of value-added IP services is increasing. These new services demand SLAs for multiple QoS classes. Bandwidth managers are capable of offering QoS at the IP layer for all transmission methods, whether transmitting over leased lines, frame relay, ATM, or even packet-over–SONET links. They offer guaranteed QoS levels by SLAs to enable service providers to gain a competitive edge in their service offering.

IP networks started out as a best-effort service with no traffic contract, which means no guarantee for delivery of data packets. With data traffic becoming mission-critical and the integration of voice and video traffic over the IP–based networks, QoS has become a major issue on the Internet and in enterprise networks. Fortunately, IP technology is readily extended to accommodate the QoS issues for insuring predictable multiservice support, and the ideal bandwidth manager's architecture extends to meet this challenge by supporting IP QoS services. These services are comprised of two critical components: the differentiated services (DiffServ) model and multiprotocol label switching (MPLS).

DiffServ provides a model for establishing multiple service levels in an IP network. It defines an architecture to provide service differentiation at edge nodes with service-provisioning and traffic-conditioning policies decoupled from forwarding behaviors. The per-hop behavior (PHB) describes the forwarding behavior, which is applied to a particular group of packets.

PHBs are implemented via the following:

- Queue management
- Scheduling
- Congestion management and avoidance

DiffServ edge nodes classify, mark, and condition packets, while DiffServ interior nodes forward packets based on priorities. DiffServ is supported in both IPv4 and IPv6 headers. DiffServ service classes are divided into the following three PHBs. They range from best effort (no guarantees) to an IP equivalent of ATM CBR (very strict guarantee). DiffServ classes are Best Effort, Assured Forward, and Expedite Forward services. Best Effort gets whatever bandwidth is left over after the other classes are processed. Assured Forward delivers a guaranteed sustained rate with bursts up to a defined maximum. The Expedite Forward delivers guaranteed bandwidth, with the least delay and packet loss. DiffServ is being widely used in most enterprise and edge data networks for delivery assurance.

MPLS is the new approach to expedite the packet forwarding while engineering the traffic, and it is a true end-to-end QoS enabler. MPLS is implemented by labeling IP packets at the ingress label-switched routers (LSRs) (see *Figure 2*). Labels are assigned as each router builds its routing table and distributed to the peer routers by means of a label distribution protocol (LDP). At the edge router, the outgoing MPLS label is determined based on source address, destination address, and priority level. The label-switched path (LSP), to traverse from upstream to downstream LSRs, is established using one of the signaling protocols (such as resource reservation protocol [RSVP] and constraint-based routed–label distribution protocol [CR–LDP]). The upstream LSR then forwards the packet to the downstream LSR. This router in turn examines the MPLS label and decides what the next hop should be. This process is repeated throughout the network until the packet reaches its destination. This approach has proven to be much more efficient, since it does not examine the IP packet headers that can cause data delay at intermediate routers. Instead, it only examines the MPLS labels that are much shorter, hence speeding up the data transfer rate remarkably.

With MPLS's ability to support different data types (including IP, frame relay, and ATM) and its ability to overlay over pure IP, ATM, and frame-relay networks, it clearly offers most promising foundation for QoS solutions for today's network. MPLS will enable the service providers to gain a competitive edge by simplifying and accelerating service provisioning, routing, and cross-connectivity. This service turn-up acceleration translates into revenues for carriers.

- Ingress LSR determines forwarding equivalence class (FEC) and assigns a label
  - Forwards Paris traffic on an LSP
  - Forwards Rome traffic on an LSP
- Traffic is label swapped at each transit LSR
- Egress LSR
  - Removes MPLS header
  - Forwards packet based on destination address

The main challenge for wide deployment of MPLS has been the lack of stable specifications. MPLS specifications are still

FIGURE 2

MPLS Forwarding Model

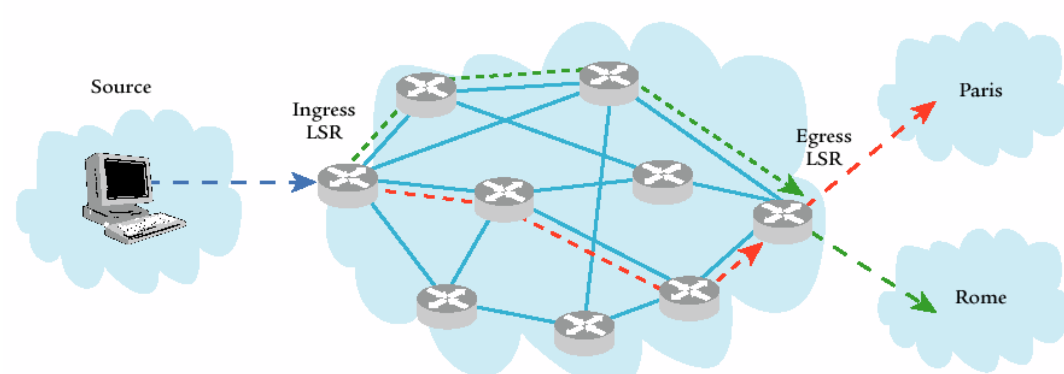

being constantly revised, and a situation has been created that does not guarantee wide interoperability among multiple vendors' products. The Internet Engineering Task Force (IETF) expects to finalize MPLS specifications by late 2001, and after that networking vendors will start shipping products with standard MPLS implementations, and then wide deployment of MPLS will occur in the network. Until then, bandwidth managers meet the IP QoS requirement by utilizing the currently defined standards such as DiffServ on platforms that will easily migrate to accommodate MPLS when the time is right.

The evolution of IP QoS from best-effort to MPLS–based services reminds us of the ATM QoS evolution and the quest for guaranteed QoS delivery on every layer of the network independent of the underlying transport technology. Bandwidth managers are strategically situating themselves in today's data market by recognizing the needs of the multi-tiered, multiservices nature of today's networks and catering to these needs.

QoS Delivery

Deploying QoS–based services is challenged by the ability to monitor and bill for such services. QoS is measured according to the SLAs between a service provider and its customers. The SLA is what requires the service provider to prove their technical competence, dedication to business, and business integrity. Today, SLAs for leased-line services are based on guaranteed bandwidth and uptime, which can be trivially supported using a dedicated TDM circuit. However, the economics of the next-generation optical network require that access and transport facilities be shared among multiple users. The ability to offer these types of services backed by SLAs requires the extensive QoS capabilities provided by ATM networks and envisioned for future MPLS–based IP networks. Due to its unique architecture that integrates TDM, ATM, and IP technologies and multiprotocol support into a single common platform, the ideal bandwidth manager provides the service guarantees necessary to support SLAs via ATM QoS today with migration to MPLS in the near future.

An essential part of a bandwidth manager is to make QoS metrics possible through smart element-management systems (EMSs). The management systems are dedicated to the true management of services and interfaces throughout the network to enable service providers to deliver accelerated, seamless end-to end solutions to their customers.

Today's technological advances allow for QoS at all layers, and multilayer bandwidth managers are designed to handle multiple traffic types while still applying QoS techniques at every layer. It seems like service providers dreams have finally come true! Now service providers, with their existing SONET infrastructures, are able to maintain the superior quality levels that they offered with SONET and to add data services to their networks while continuing to be profitable. A protocol-agile bandwidth manager fits perfectly into both the present network and the next-generation network, supporting multiple service classes and maintaining the highest QoS. The time is here to take advantage of QoS your way!

Distributed Fault-Tolerant/ High-Availability Systems

Harry Singh

Product Manager

Trillium Digital Systems, Inc.

Executive Summary

Today's telecommunication platforms must provide "five-nines" availability, meaning practically no loss of service due to hardware or software errors, nor any downtime for software upgrades or hardware maintenance. This expectation places an unprecedented burden on service providers: to ensure that all of the network elements needed to support a service are functioning whenever a user requests that particular service.

Five nines means 99.999 percent availability. Achieving 100 percent perfection, while desirable, may be impossible, and a fault-tolerant system is designed to function correctly given two predictable contingencies: one, the small (0.001 percent chance) but practically inevitable presence of a fault in the system, and two, the likely requirement for periodic maintenance or system improvements. We will return to this point.

Supporting five-nines service availability depends on the near-flawless interaction of software (application, operating systems (OS), and management middleware), hardware and network design, as well as environmental and operational factors. Hardware fault tolerance typically relies on redundant processors, memory, buses, power supplies, and databases. Software fault tolerance uses a combination of software redundancy and simple hardware redundancy to provide the necessary availability in the case of failure. Network fault tolerance uses redundant data link cross-connects (T1s, digital signal [DS]–3s, optical carrier [OC]–3s, etc.).

This paper gives an overview of fault-tolerant/high-availability (FT/HA) components. *Figure 1* shows where all the FT/HA components of a system can reside. Both the hardware and software solutions and the challenges addressed by these solutions are discussed. This paper also outlines Trillium's FT/HA and the patent-pending distributed fault-tolerant/high-availability (DFT/HA) architectures and implementations, and how they are designed to overcome these limitations.

FT/HA Concepts

High Availability

To understand availability, we first need to understand reliability. The reliability of an element is the conditional probability that the element will operate during a specified period of time. A system may be considered highly *reliable* (that is, it may fail very infrequently), but, if it is out of service for a significant period of time as a result of a failure, it will not be considered highly *available*. One measure of an element's reliability is its failure rate, or mean time to failure (MTTF), which is the interval in which the system or element can provide service without failure. Another measure of reliability is the mean time to repair (MTTR), which represents the time it takes to resume service after a failure has been experienced. Systems may go out of service for any number of reasons, such as the occurrence of a fault, repair activities, software loading, hardware upgrading, or periodic maintenance. In any such case, MTTF and MTTR immediately become vitally important.

The availability of an element is the probability that the element is in service and available to a user at any instant in time. It can be expressed using these measures of reliability.

$$Availability = MTTF/(MTTF+MTTR)$$

As the equation shows, the availability of systems can be increased by designing components that are highly reliable (high MTTF) and/or by shortening the time required to repair the system and return it to service (low MTTR). Since it is impossible to create systems that never fail, the key to high availability is to make recovery time as brief as possible.

Network elements typically operate with a target of five-nines availability. Such a level of availability in a telephone switch, for example, means that the switch is expected to be out of service only for about five minutes *per year*. These few minutes are all the time needed to repair faults, load software, upgrade hardware, and perform periodic maintenance and any other necessary activities.

FIGURE 1

FT/HA Components within a System

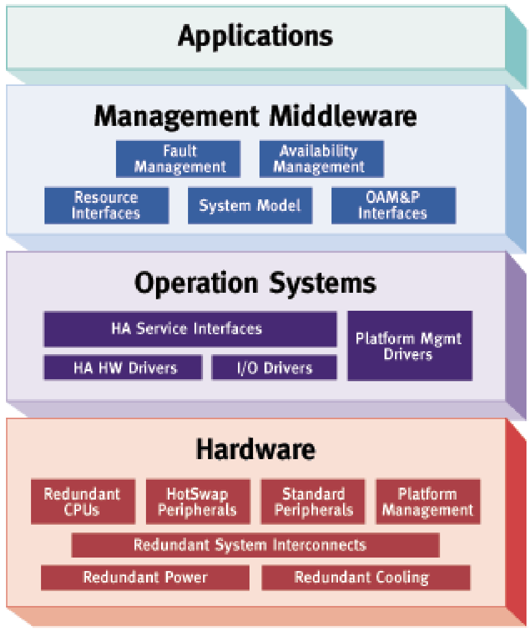

Fault Tolerance

A fault-tolerant system is available in the presence of faults. The design, development, and testing of highly available systems is challenging and expensive. It generally proves more economical to assure high availability by detecting faults and avoiding service disruptions through redundancy in the system, that is, designing fault tolerance into the system so that it functions correctly and is continuously available in the presence of faults. Fault tolerance will increase the availability of a system, which is measured by the probability that the system will operate and be accessible when required for use, even during periods of preventative maintenance or repair.

The simplest failure-response strategy is to let a non-redundant system fail and then repair it off-line. This is the lowest-availability strategy, and it is unacceptable for mission-critical computer and communications applications in which high availability is a necessity. Such applications require systems that include fault-tolerant capabilities that make them highly available.

Fault tolerance can be achieved through a combination of system hardware and software. One way to achieve system fault tolerance is to use redundant hardware components within the system (e.g., multiple processors, buses, or power supplies in a single system) that operate simultaneously and in parallel, comparing results of the operations performed. This is referred to as a hardware fault-tolerant or redundant system. Using software to handle the fault management of two or more subsystems, regardless of whether each subsystem has redundant hardware components, can enhance system fault tolerance. When one subsystem fails due to the presence of a fault, the other subsystem takes over, so that the overall system can continue to operate without disruption in service.

Periodically causing a switchover between the software modules and restarting the passive software module can achieve software redundancy.

As we mentioned in defining "five nines," failure is just one of the contingencies that must be overcome to achieve 99.999 percent availability. The other and more likely event is anticipated maintenance to the hardware or software. Fault tolerance is necessary to enable the system manager to plan and execute "rolling" upgrades, that is, controlled, programmed improvements that avoid system shutdown. In a fault-tolerant system, five-nines availability can be ensured because the system manager can execute a rolling upgrade at an optimal moment—an anticipated fix with no downtime instead of a remedial one with a minimized interruption.

Market and Technology Drivers

Following are some of the market drivers for FT/HA systems.

Converged and Decomposed Network

Telephone companies and their equipment manufacturers have set standards for five-nines availability that other network elements in a converged and decomposed network must meet. Internet infrastructure and services will need to strive for near-perfection to provide services seamlessly across both the Internet and the public switched telephone network (PSTN).

These decomposed network elements (signaling gateways, media gateways, media gateway controllers, etc.) are built using commercial off-the-shelf (COTS) hardware and software solutions. Because of open standards and architecture, these solutions are obtainable from multiple vendors, and not, as before, from a single vendor that supplied everything—applications, hardware and software—in one box. These COTS solutions must meet the stringent availability expectations set by the PSTN.

Internet

The growing popularity of the Internet as a channel for business transactions and as an alternative to the traditional voice network is driving manufacturers to provide equally reliable and near-continuously available systems. Internet customers expect to be always on and always connected (AOAC), and they demand quality and reliability levels such as those established by the telecom industry during the past 10 to 15 years in the PSTN. Convergence is driving these same requirements into data networks and mainstream infrastructure equipment.

In addition to mission-critical data-processing servers and carrier switches, Web servers, signaling gateways, media gateways, media gateway controllers, wireless base-station controllers (BSCs)/mobile switching centers (MSCs), and computer telephony (CT) servers are expected to have higher levels of availability. *Figure 2* shows many of these elements in a decomposed network.

FT/HA Hardware Solutions

Designers and developers can take several design approaches to ensure high availability for individual network elements. The simplest type of network element is non-redundant and must be repaired off-line if it fails. This

FIGURE 2

A Decomposed Next-Generation Network

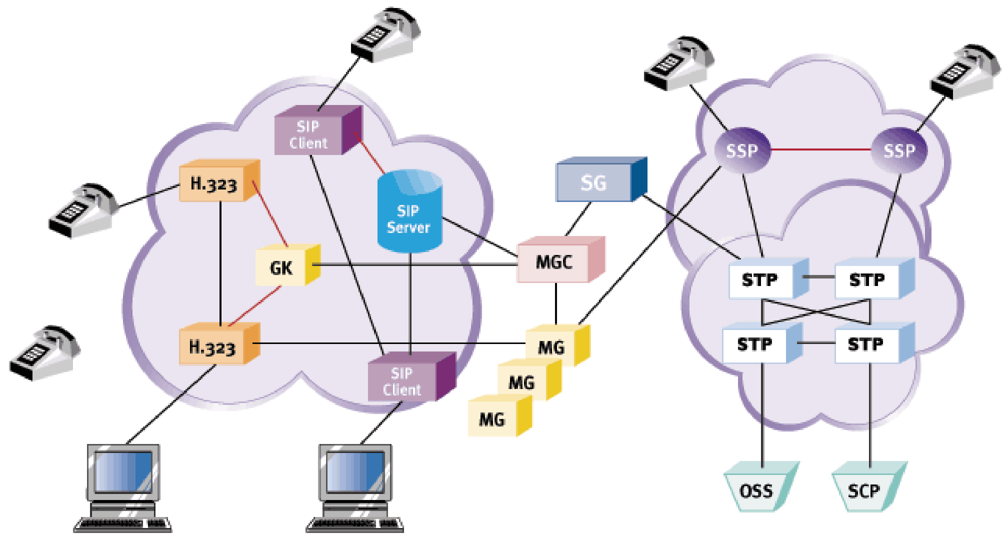

type of element will have relatively low design complexity and low cost. (Depending on its reliability, it may also have low availability.) In contrast, HA elements require both the ability to support on-line repair (usually through the "hot swap" of components while the element is in service) and additional redundancy. Hardware fault tolerance typically relies on redundant processors, memory, buses, bus cross-connects, power supplies, cooling systems, and disk storage. *Figure 3* shows two extremes of these solutions.

Elements with additional redundancy typically use "retry" and "masking" for recovery. Retry-based elements attempt to ensure that there is a second attempt at the operation if an initial operation fails. If the second attempt succeeds, the fault was probably transient. If the second attempt fails, the fault is probably permanent. Masking-based elements attempt to ensure that only the results from the correct operating portion of the element are used if a component fails. In either case, if a component has failed, the system attempts to detect, diagnose, isolate, recover, repair, and compensate for the fault.

System designers can enhance the fault tolerance of a system by combining simple hardware redundancy with fault-management software: replicating hardware subsystems to provide the necessary backup hardware in the case of failure and executing a copy of the same software on each node provide simple hardware redundancy. These hardware subsystems, or "nodes," can either be entire computers complete with processor, memory, buses, power supplies, cooling systems, and disk storage (in practice this is called "clustering") or any of these components individually. The fault-management software manages the switchover from the failed node to another operational node, maintaining the node-state information of the failed node so that it can be used by the system after the switchover. This eliminates the need for complex circuits, typical of more complicated hardware fault-tolerant approaches, and can significantly decrease the design complexity and cost of the system.

In telecommunications and networking applications, system designers typically achieve fault tolerance by using

FIGURE 3

Spectrum of FT/HA Hardware Solutions

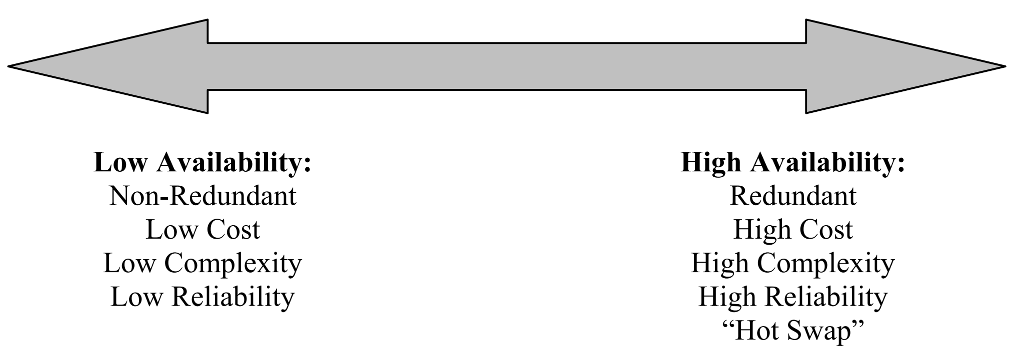

either dual-node or multi-node architecture. In a dual-node architecture, two nodes are running simultaneously. In some systems, both nodes are active. In other systems, one node is assigned to be active and the other node is a standby. A more scalable multi-node architecture allows for the number of active and standby nodes to be configured and for the active nodes to share the system load. The dual-node architecture is more popular because of its lower cost.

A dual-node system typically achieves fault tolerance in one of two ways:

- Both nodes are active, sharing the load of the system by executing different tasks. If one active node fails, the other active node takes over the tasks of the failed node.
- Both nodes are capable of executing the same tasks, but one node is active and the other is standby. If the active node fails, the standby node becomes active and takes over the assigned tasks.

In either case, the switchover from one node to another is either controlled (operator-initiated) or forced (system-initiated). Generally, an operator-initiated switchover is performed for maintenance purposes, whereas a system-initiated switchover is carried out when a node fails. Application software is usually unaware of underlying HA hardware. These hardware nodes communicate with each other via current technologies such as cPCI buses or some upcoming technologies such as InfiniBand, RapidIO, or switched Ethernet.

Platform-management software provides the monitoring and controlling capabilities required for the detection, diagnosis, isolation, recovery, and repair of nodes.

Three common redundancy schemes found in hardware-based solutions are as follows:

- *2N:* One node in standby for every node in operation. For example, two independent cPCI chassis providing redundancy at the central processing unit (CPU) and input/output (I/O) level.
- *N+1:* One standby node for N operational nodes. For example, two cPCI chassis with crossover providing redundancy at the CPU and I/O level.
- *N+M:* A pool of N nodes working in normal operation with a pool of M nodes in standby. This is more typical of distributed systems, where nodes are active and also the standby of each other.

FT/HA Software Solutions

As discussed before, hardware fault tolerance usually relies on redundant processors, memory, buses, power supplies, and disk storage. Software fault tolerance, on the other hand, uses a combination of software redundancy and simple hardware redundancy to provide the necessary availability in the case of failure. Unlike hardware fault tolerance, software fault tolerance is absolutely necessary to provide system-wide redundancy. A major technology differentiation of HA systems is the inclusion of configuration and fault-management software functions, enabling the detection, diagnosis, isolation, recovery, and repair of faults anywhere that could result in the failure of the HA system. This section discusses how the OS, management middleware (MM), and application software can provide FT/HA functionality.

Operating System

HA systems put certain requirements on the OS. Some of the important requirements are as follows:

- *Memory Address Space Protection for Kernel and Applications:* Support for virtual address space. With this type of memory protection, the active process cannot corrupt the memory space of other kernel processes and applications in the system.
- *Robust Kernel:* The OS should be able to isolate and prevent the propagation of, or mask the impact of, potential hardware and software faults.
- *Ability to Deal with a Dynamic Hardware Configuration:* An HA–aware OS must be able to support the dynamic reconfiguration of applications and drivers as faulty hardware is replaced with healthy hardware. This "hot swap" support should include fast restarting of applications and dynamic installation, and an update of device drivers.
- *Support for Fail over Mechanisms:* The OS should be able to respond rapidly to faults and fault-recovery procedures. The OS should be able to support fault recovery as dictated by policy management.
- *Interface to MM:* For the management of the OS itself and the ability to report faults externally.
- *Support for On-Line Upgrades*

Management Middleware

Management middleware provides a set of configuration and fault-management capabilities. MM typically contains direct interfaces to the OS, to hardware devices, and to the applications, and it adds additional capabilities, such as messaging or state replication services, to an OS.

The following are some other important capabilities of MM:

- *Fault Management:* Detects, diagnoses, isolates, recovers, and assigns the role and repairs the faults anywhere in the HA system. MM may also support prediction of faults by looking at changes taking place in the system.
- *Availability Management:* Collects data, does role assignment, and performs rapid recovery in conjunction with fault management.
- *Resource Interfaces:* Provides interfaces to the OS and application.
- *Systems Model:* Manages configuration and dependencies among system-wide components. This also maintains a state-aware model of the total system.
- *Operations, Administration, Maintenance, and Provisioning (OAM&P) Interfaces:* Provides local and remote interfaces to view and manage the system or its components for control and intervention.
- *Messaging Services:* Supports checkpointing and inter-application messaging.
- *Support for On-Line Upgrades*

Application

The application software may or may not be aware of the underlying HA system infrastructure. Application software will typically interact with the OS and MM no differently than in non–HA systems. In case the application needs to be made aware of the underlying HA system infrastructure, there must be an interface between the application software and the MM. This interface may provide access to services unique to the HA system, such as checkpointing application state or heartbeating.

Current Challenges

Legacy fault-tolerant systems were built using special fault-tolerant hardware platforms that were, in many cases, designed specifically for the application/service that they were to host. In addition, fault-tolerant applications were constructed from scratch and were expensive to design, produce, and maintain. Along with specific hardware and applications, the system-management software was also designed for the specific system that it was supposed to manage. In such cases, the FT/HA software might be available in binary form only, rather than source form. This would limit the available platform choices. For the system designer to be able to select a different board, processor and OS for a specific application, the availability of portable software—written in a popular programming language—is paramount. Otherwise, system designers and engineers will be forced to make a platform choice that may not be suited to their needs: They will be precluded from choosing alternative platforms that would work better for their application or provide cost or other competitive advantages. Moreover, they will not be able to modify the software to provide their own proprietary value to the end product.

Some other important limitations of current software solutions include the following:

- They execute the same software simultaneously on both active and standby nodes. Using this solution, the entire system could fail, since the same fault could occur in the software on both nodes at the same time.
- Some systems are unable to protect individual software instances separately if both the instances are on the same node (either active or standby). This amounts to an "all or nothing" solution. A more flexible approach is to distribute the protected instances on separate nodes as appropriate for their application. The benefit of this approach is easier fault detection.
- Some systems are unable to detect and resolve sequential faults. It is possible for an error to occur during the recovery period of another fault. The second error must also be considered and resolved in a graceful manner (e.g., first queued, and then processed after the first error has been resolved).
- Some systems are incapable of rolling upgrades and therefore are to be taken off-line to perform planned maintenance.

Trillium's FT/HA Solutions

Trillium's source-code software solution provides a flexible framework that is platform-independent and cost-effective. Trillium provides FT/HA communications layers and stacks that can be classified as applications that provide service to a system user. This FT/HA solution maintains active calls during software and hardware failures. Provided in source code, the Trillium solution gives systems designers and engineers total freedom when choosing their hardware platform and OS (kernel or user space). Developed primarily for telecommunications products, the Trillium solution is also applicable to other types of products requiring high availability.

Trillium's solution does not employ the standby node to execute the protected application at the same time as the active node, thereby protecting the system from "mirror-image" errors that could occur on both nodes. Instead, the active node updates the standby node's internal states, so that when the standby node becomes "active," it will start executing the software based on the most recent update. This approach does not allow the propagation of a software fault from the active node to the standby, since the standby is not executing the same segment of the code at the same time as the active node. Therefore, the system can potentially bypass a software problem that may occur in the active node.

Trillium's solution includes an FT/HA core product, plus a series of technology-specific products, each called a FT/HA protocol-specific function (PSF). Trillium's FT/HA core software product can be used on systems consisting of either hardware–fault-tolerant nodes or non–hardware–fault-tolerant nodes.

The FT/HA core enables system designers and engineers to replicate a node and turn those nodes into a FT/HA system. Trillium's solution can be used for open service platforms, which in turn can be used as a basis to create other products. For example, it can be used for a service control point (SCP), which accesses databases for telephone number translations. The Trillium-enabled SCP is able to use the dual-node architecture using either a telecommunications board manufactured by any of the major hardware providers or a board developed in-house, and it is able to use any multi-tasking OS. For other systems, such as an intelligent peripheral (IP), the multi-node approach using the Trillium solution can provide the high level of availability that is required by such systems.

For specific telecommunications applications, Trillium develops the FT/HA PSF for particular technologies (e.g., signaling system 7 [SS7], asynchronous transfer mode [ATM], integrated services digital network [ISDN], H.323), providing a complete solution for FT/HA systems. The user need only place an application, such as a mobile application part (MAP), on top of this stack and configure the protocol stack for its particular requirement specifications.

Trillium's FT/HA solution provides the following functionality:

- Maintenance of all active calls during software or hardware failures by updating the states of the standby node without execution of its state machine.
- Alarming the stack manager software when a fault occurs within the protected layer and carrying out a forced switchover. The stack manager is a system-dependent function and must be implemented by the system designer. Once the stack manager has received a failure indication, it can initiate corrective action by invoking Trillium's system manager to carry out a forced switchover.
- Protection of a single application (layer) or multiple applications (stack) with each layer residing in separate nodes or multiple nodes. When a node fails, Trillium's

FT/HA software will reroute the messages destined for the failed node to another operational node.

- Trillium's FT/HA software will work in a system in which the stack manager carries out a controlled switchover for operator-initiated routine maintenance. The stack manager "requests" the switchover, and the software subsequently carries out the "request."
- Trillium's FT/HA software will work in a system designed to detect and resolve sequential errors. The stack manager must queue the errors and notify the software.
- Trillium's FT/HA software supports rolling upgrades, thus ensuring five nines of availability even during planned maintenance.

Most equipment manufacturers use a dual-processor architecture with active and standby subsystems. Trillium's FT/HA solution enables the standby subsystem to maintain state information through state updates from the active subsystem. These updates prevent the loss of state information, enabling an orderly switchover procedure from the failed subsystem to the standby subsystem. *Figure 4* illustrates this functionality.

Trillium's DFT/HA Solutions

Trillium's portable FT/HA solution has been extended to include distributing the processing load across multiple processors. These solutions, called DFT/HA, provide the high performance and scalability demanded by network infrastructure manufacturers.

The patent-pending DFT/HA software, coupled with Trillium's broad suite of communications software solutions, will enable manufacturers to build products meeting today's stringent carrier-grade telecommunications requirements. The Trillium's DFT/HA products are designed for the active/standby dual-node architecture and scaleable multi-node architecture, which allows for the number of active and standby nodes to be configured and for the active nodes to share the system load. In this way, different processors can be active and standby of each other, providing an N+M configuration. *Figure 5* shows this functionality. This N+M configuration goes beyond the pure distributed solution, with no standby nodes in the system. In the non-redundant solution, if the active node goes down, the services provided by that node are no longer available.

Trillium's solution includes the DFT/HA core product and a series of technology-specific products, each called a FT/HA PSF and load-distribution function (LDF). The LDF distributes the traffic of a protocol layer onto multiple hardware nodes to increase the traffic-handling capacity of the protocol layer. Currently, Trillium's DFT/HA solutions include the popular transactional capabilities application part (TCAP) (TCAP, SCCP, and MTP3) and ISDN user part (ISUP) (ISUP, SCCP, and MTP3) stacks.

These solutions add the following functionality to the existing FT/HA solutions:

- *Flexible Architecture:* DFT/HA allows manufacturers to design distributed FT applications, pure distributed applications, and pure FT applications. A DFT/HA solution supports 2N, N+1, and N+M (same as N+1, when M is 1) redundancy schemes.
- *Load Distribution and Scalability:* DFT/HA supports multiple active processors simultaneously with multiple

FIGURE 4

Trillium's FT/HA Solutions

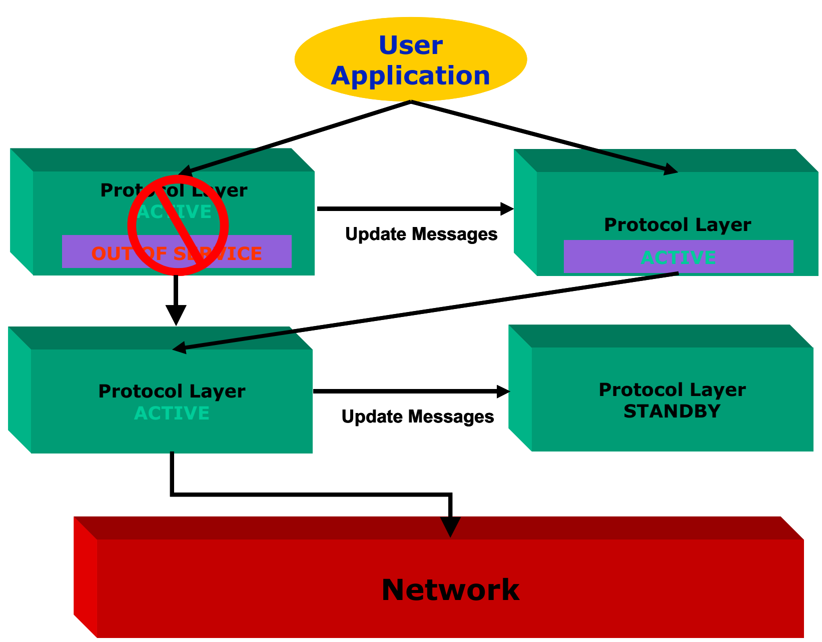

FIGURE 5

Trillium's DFT/HA Solution

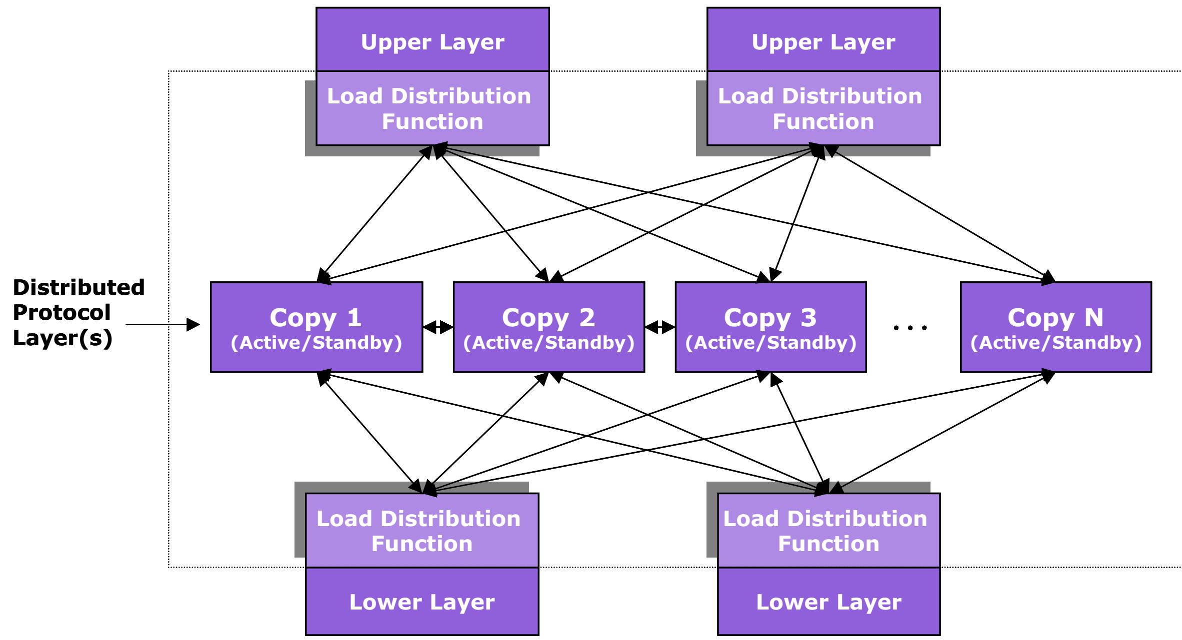

standby processors, thus avoiding the need to reconfigure the system's architecture when it needs to be expanded to handle the extra traffic.

- *Smooth Migration Path:* DFT/HA allows the co-existence of both distributed and non-distributed software layers within a protocol stack.

Summary

HA products and services will play an increasingly important role in the network of the future. The availability of services is a complicated function of the equipment design and also of the network design. Different design choices will result in different levels of complexity, cost, performance, potential information loss, and (of course) availability. The reliability and availability of the public telephone network has set the standard against which future services from an integrated Internet and PSTN infrastructure will be measured.

For mission-critical applications such as telephony, FT/HA systems are vital. To achieve FT/HA, Trillium has undertaken a substantial development effort in recent years to enable it to offer a FT/HA solution that is flexible and platform-independent. Because it is offered in portable source form, Trillium's FT/HA software solution does not force users to limit their choice of hardware platform and OS (unlike binary solutions). Trillium's solution gives system designers and engineers the freedom to select the optimal platform for their specific application and to focus on ways in which they can add proprietary value to their products with an optimized investment of time. In addition to supporting Trillium's existing fault-tolerant capabilities, the new patent-pending DFT/HA architecture enables high performance and scalability in a system by distributing a protocol layer across multiple processors while maintaining fault-tolerance support.

OSS Microcosm: Time to Think Afresh

Ashish Srivastava
Practice Head, OSS
Wipro Technologies

Dinesh Rathi
Consultant, Process Management
Wipro Technologies

Introduction

We have all seen the network building frenzy that caught on like a bug last year. Operations support system (OSS) implementation followed the trend. The best and only the best was the *mantra*. What ensued is now a well-known fact. A rare few exist to share the experience of those mindless implementations with little strategy. It is rightly so that today the service providers are ever more cautious to safeguard their investment and make sure that every penny of OSS worth is brought in at the right time at the right price.

It will be, however, too simplistic to assume that everything that is happening on the service-provider front is right from the OSS standpoint. Rapid changes in the business plans of organizations keep making the OSS solutions obsolete, even before they are implemented. The organizations are creating and recreating themselves almost on a monthly basis in an attempt to discover the magical formula for success. The ideal product mix is still a mirage. These are definitely no ordinary times. What is required is a fresh approach to the OSS implementation.

Pitfalls of Classical OSS Strategy

The classical OSS strategy is based on many assumptions, which need revalidation in the current context. Some of them are being covered here.

Myth 1: Organizational Model Needs a Focus
How the service providers started being "network-centric" has long been discussed. They then became "service-centric," graduating to "customer-centric" model. Such hair-splitting discussions have lost their value in the current scenario. There cannot be one center of this complex world of OSS at any point of time. Try to focus on the customer, and you can safely bet that your internal organization will never catch up with the customer demands. Focus inwardly and be sure that the customers have walked away before you can say "VPN." Start debating about your product mix, and the yawning customers would have moved on while your sales department waited on the sidelines like an army without arsenal.

How does the OSS strike the right balance?

Myth 2: OSS Is a Strategically Well Planned Affair
The Big Bang approach has long since been rejected. The recommendation was to iteratively evolve the OSS in a well-planned manner. So where is the catch. Well, the fact is that there is too much flux in the operational plans of the service providers. There is no long-term basis to act as a foundation for the long-term OSS strategy. The OSS solution has to contend with five cardinal truths of today's business:

- Products will change.
- Market will change.
- Processes will change.
- Organizations will be ever evolving.
- Solutions will be ever evolving.

Where in the world can we find a super-flexible OSS that allows adapting to the needs of the service provider at will? Even the loosely coupled OSS integrated through middleware architecture can go haywire in the first few cycles of change.

Looks like we have hit the wall when it comes to OSS implementation.

Myth 3: OSS Comprises of Independent Functional Layers
The TMN Forum's functional hierarchy of layers of OSS management is etched in every mind related to the field of OSS. It is an undisputed fact that this structure by itself brought a lot of clarity to understand the OSS implementation. However, this model leaves a lot of holes in terms of adequately relating the OSS to the rest of organization and the customer.

OSS is no longer a set of functional layers that can be managed independently. The interrelationship of OSSs and business support systems (BSSs) runs deep within the organizational processes. Unless a bond between the business and the solution is established, the chances of success are dim.

OSS Microcosm

It is time to look at OSS in a different way. These are times that require reliability, flexibility, and economy of scales all together. To achieve this, we have to break the old-world view of an OSS being just a bundle of individual or integrated solutions. The OSS has to be viewed as an integrated mesh of varied factors. Change one of them, and you have a ripple effect on the others. It is important to move all of them together in any evolutionary model of development. We will call this interrelationship an OSS microcosm.

Figure 1 shows the visual representation of an OSS microcosm.

FIGURE 1

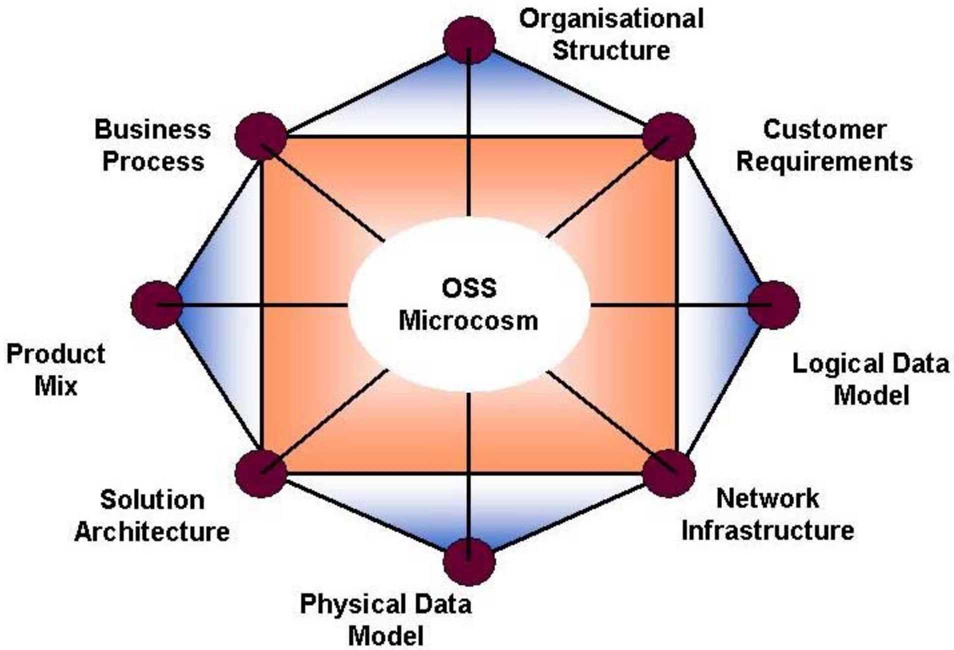

Microcosm Management

The very complexity of the model may dissuade the systems integrators (SIs) as well as service providers with the enormity of the task of managing and implementation. The problem, however, is not such a discouraging one. All it requires is a combination of methodology and discipline. The following three simple tenets, if followed to the core, make the whole issue a manageable one (see *Figure 2*):

- Manage business processes.
- Manage traceability.
- Manage change.

Let us look at each one of them, one at a time.

Business-Process Management

The business processes are often treated in isolation. They are often treated as independent to the core OSS solution. Business wisdom dictates that the business dictates the process and that the process dictates the solution.

However, we often fall into the trap of technology and follow the aforementioned principle in reverse by allowing the solution to dictate the process. This should be watched for carefully. In the short run, the application-dictated process leads to fast implementation but proves to be inadequate for the growing business needs.

The key to efficient process management lies in the way the processes are modeled and represented. A hierarchical approach of process representation is the most common one and serves the purpose well. A typical hierarchy includes the following:

- *Level-1 Processes*: Indicating the functional organization.
- *Level-2 Processes*: These are the processes at the organizational level.
- *Level-3 Processes*: Represent the break-up of the level-2 processes to the next level of detail. These include the sub-activities, the input and the output details, and the business/validation rules.
- *Level-4 Processes*: These are the operational procedures used by the end users.

Using a tool for process modeling is a good idea. Tools allow for greater accuracy, version control, and advanced features such as simulation to check out the efficiency of the process.

Some important points to note in process modeling are as follows:

- All of the processes do not get automated on day one. It helps to segregate the manual and automated processes (preferably at level 3). This helps in keeping track of the changes to operational procedures as the solution evolves and encompasses more and more of the process.
- The process handoffs should be carefully noted so that the integration can be carefully tracked.
- The level-3 processes should be at a level of detail at which they can be mapped to the solution requirement.

Traceability Management

We have discussed before the interdependencies between various entities of the microcosm. Without a formal way to track the traceability between these entities, it will be but a matter of time before chaos takes over. There will be no way to judge the impact on any one of them due to the others.

FIGURE 2

However, traceability management is no joke. From the OSS perspective, I recommend that traceability to be managed within the quadrant in *Figure 3*.

It is very important to note that the business processes translate into a solution through two different paths:

- Through the functional requirement
- Through the data model

It is also necessary to ensure that the traceability within these entities is maintained for all of the products. It is true that a large part of the solution or processes may not change with products, but then there will be some product-specific parts that need to be tracked as well. What works well for one product may not always work well for another.

Business processes at their detailed level are nothing but functions and data. Functions are important for the end users. They understand the language of functionality that they would need to run their part of the business. This functionality is what they would expect to see in the solution.

A solution architect's point of view is, however, different. It is imperative to have a robust architecture that is scaleable and devoid of data redundancy and data integrity issues. This is achieved by mapping processes to a data model, which should then translate into the solution architecture.

Unfortunately, there is no single tool that can claim to manage all of this traceability together. Lack of a single tool often discourages integrators to map traceability efficiently. The result is a loss of the holistic picture and a divergence of the processes and the solution. However, the tool issue is often exaggerated. There are off-the shelf tools that help to map the processes to the data model. Then there are tools that derive the data model from the solution. There are a host of requirement-management tools that help to keep traceability a requirement of the solution. A good combination of these will achieve the purpose. If managing multiple tools is an issue, the good old Excel spreadsheet is better than having no traceability at all.

FIGURE 3

OSS Microcosm: The Traceability Quadrant

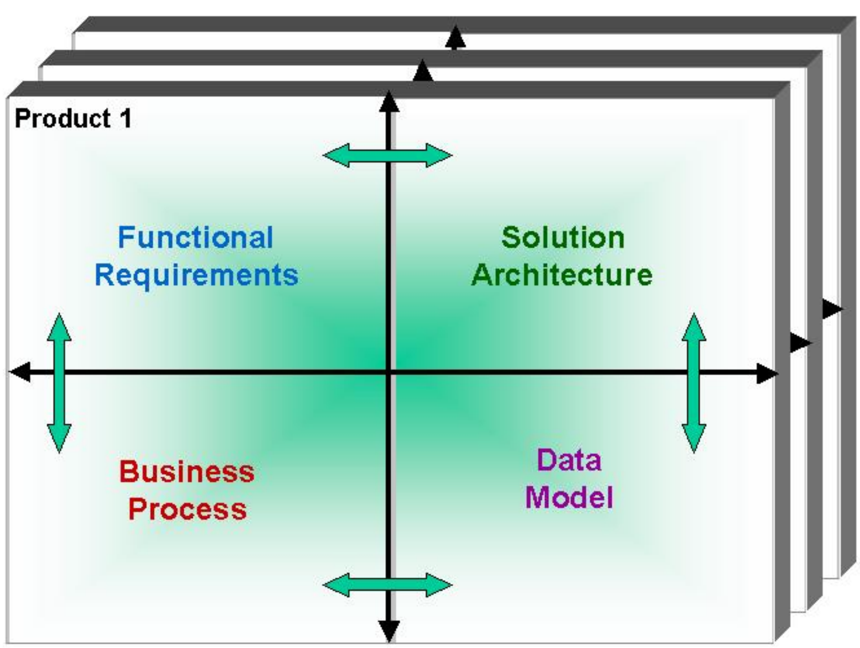

Change Management

The last important tenet for managing the OSS microcosm is having a controlled change management. We touched upon the inevitability of change at the beginning of this paper. The issue is to manage that change so that if one entity changes, the impact on all of the other entities is taken care of. This ensures that the interrelationship retains its integrity.

From the OSS perspective, change can occur from many sources. Some common change triggers are as follows:

- Organizational structure change
- Business-process change
- Functional requirement change
- Solution changes

If the previously discussed issues of process management and traceability have been well addressed, the change management is a cakewalk. The crucial steps in change management are as follows:

- Initial assessment of the change
- Notification of the change to all stakeholders
- Impact analysis of all the affected entities
- Change implementation strategy

It should be noted that it is not always possible to propagate changes to all of the affected entities instantaneously. There will be a period of time when the discrepancy in the interrelationship has to be managed. This requires a clever strategy.

Figure 4 shows an example of how an interim discrepancy can be managed through a temporary stage of work around. At the earliest next opportunity, all the entities are brought in synch and new traceability is established.

Summary

In the current times, it is important to see the OSS as being an integral part of a meshed relationship in a service provider's business. Most important are the relationships of the OSS solution to the business process, functional requirement, data model, and product offering. We have termed this interrelationship as an OSS microcosm.

To be able to manage the microcosm, it is important to manage the business processes, traceability, and change control. With the aforementioned approach in place, OSS is likely to be in sync with the business objectives and will be able to sustain its integrity in a rapidly changing environment.

FIGURE 1

OSS Microcosm: The Change Life-Cycle

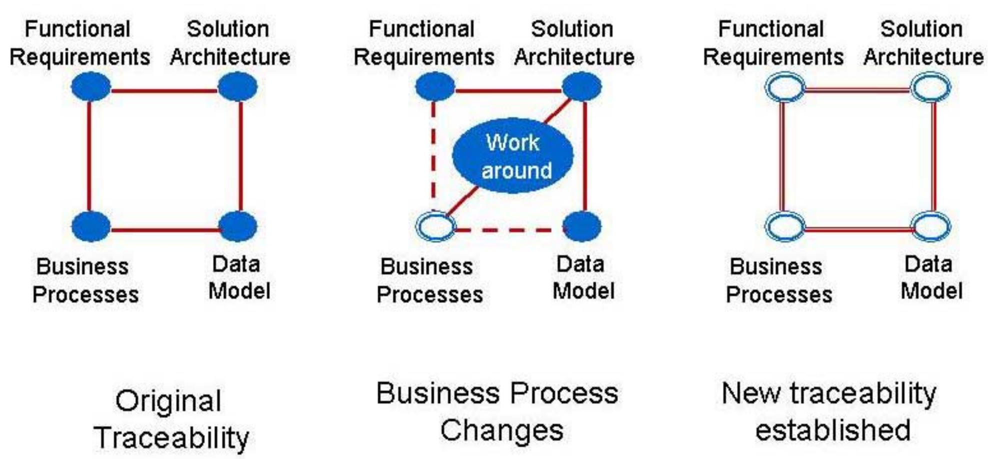

Real-Time Billing in SIP

Baruch Sterman

Chief Scientist

deltathree

1. Abstract

Session initiation protocol (SIP) [1] is fast becoming the protocol of choice for Internet protocol (IP)–based communications—specifically telephony (voice over IP, or VoIP), video, and instant messaging. SIP's strengths lie in its e-mail–like addressing scheme, its plain-text headers (similar to hypertext transfer protocol [HTTP]), its scalable architecture, and most of all in its ability to carry generic data (following the multipurpose Internet mail extensions [MIME] standard of e-mail fame). Though interconnecting a SIP call leg with the legacy phone system (public switched telephone network, or PSTN) is well defined and communication flows seamlessly, the billing model regnant in the telephony world does not mesh quite so naturally with SIP. For example, the ubiquitous pre-paid calling-card services require querying the account status before a call is authorized and a real-time adjustment of the balance as soon as a call terminates. SIP does not address either of these functionalities. This paper will examine the requirements of such systems and suggest possible implementations within an SIP architecture.

2. Review of SIP

2.1. Introduction to SIP

SIP is a signaling and session-management protocol that is meant to establish a communication session between two IP–based endpoints. SIP provides the capabilities to do the following [2]:

- *Find the Target Endpoint:* Endpoints are addressed in a format similar to e-mail, i.e., user@InternetAddress, where the Internet address may be either an IP address or a fully qualified domain name, such as "deltathree.com."
- *Negotiate a Common Method of Communication:* Different users may have different capabilities. In a SIP message, each user lists its own set of codecs that it supports and the precedence of each. This information is carried as session description protocol (SDP) in the body of a SIP message.
- *Manage the Session:* SIP messages are defined to set up and tear down a communication session and can pass information between the two endpoints outside the media channel.

2.2. SIP Network Elements

The basic elements that form the building blocks of a SIP network are the following:

- *User Agent:* These are the endpoints that initiate and terminate calls. Examples of user agents are IP phones, software clients, and IP voice gateways that interconnect with the PSTN (the legacy telephone network).
- *Proxy Server:* Proxies forward SIP messages on to other SIP network elements according to rules that are defined within the proxy.
- *Registrar:* The registrar is used to dynamically track the location of user agents and forward messages to their provisional address. Typically, a user agent will "register" with a registrar and provide the information regarding its current address (i.e., IP address). The registrar will then forward calls destined for the registered user on to the specified location.
- *Redirect Server:* This server is similar to a proxy, but instead of forwarding messages, it responds to SIP requests with the location of a user agent or server. The entity that sent the request will issue another request directly to the location specified by the redirect server.

Often, proxy software contains proxy, registrar, and redirect functionality. A proxy may also keep a set of rules that determine how to route calls when the attempt to establish a connection with a user at its default location fails. An example of such "location service" rules would be to reroute a call to another SIP address when the primary address is busy or unavailable.

2.3. SIP Messages

The common SIP messages are the following:

- INVITE: Initiated by a user agent inviting a second user agent to participate in a session.
- ACK: Acknowledgement that is returned confirming that the user agent has received a final response (as opposed to a provisional response, such as "trying" or "ringing").
- BYE: Indicates that one endpoint user agent is terminating the session.
- REGISTER: Contains the location information necessary to reach a user agent.
- CANCEL: Indicates that a pending request should be cancelled.

2.4. SIP Responses

SIP responses are similar to HTTP. Responses are numbered, with the first digit indicating the type of response. *Table 1* shows the different response types.

TABLE 1

SIP Response Messages

Response Group	Description
1XX	*Informational:* Indicates status of call prior to completion.
2XX	*Success:* Request has succeeded.
3XX	*Redirection:* Server has returned possible locations.
4XX	*Client Error:* The request has failed due to an error by the client.
5XX	*Server Failure:* The request has failed due to an error by the server.
6XX	*Global Failure:* The request has failed.

2.5. Why SIP

SIP offers the following benefits:

- *Simplicity:* SIP messages are encoded in plain text and are easy to understand and troubleshoot. A SIP stack is much lighter than other VoIP stacks.
- *Distributed Functionality:* The proxy/forward approach utilized in SIP encourages distributing intelligence throughout the network. A call may flow from proxy to proxy, each working on the call with its own logic and functionality. Changes can be made to one component while leaving the others untouched.
- *Extensibility:* SIP is very flexible in that the message body can contain attachments similar to e-mail. User agents are free to use SIP for any type of communication, as long as both endpoints can process the attached information.

2.6. SIP: Past, Present, and Future

SIP was initially proposed as a method of inviting users to join in large-scale multicast conferences. It was soon realized that SIP could be used to establish point-to-point IP phone calls as well. By 1999, SIP was a proposed standard for establishing communications sessions. Over the next few years, device manufacturers, software vendors, and service providers got together at "bake-offs" to test the interoperability of their SIP components. In November 2000, the third-generation partnership project (3GPP) wireless initiative adopted SIP as its protocol. Further validation of SIP came in the spring of 2001, when Microsoft announced that its Messenger platform and voice services (released in the fall of 2001) would be SIP–based. Device manufacturers and application developers are continuing to roll out SIP–based products and services.

3. SIP Call Flows

3.1. A Simple SIP Call

To get a better understanding of SIP and how it works, an example of a call flow between a user agent, a proxy server, and a voice gateway is described (see *Figure 1*).

The user agent initiates a call by sending an INVITE message to the proxy server. The INVITE message contains the information necessary for the target endpoint to open a media channel to the user agent. This information includes the IP address and port of the user agent, as well as the codecs that the user agent supports. The proxy forwards the INVITE message to the voice gateway. (When the proxy receives the INVITE message from the user agent, and when the voice gateway receives the INVITE from the proxy, a 100 Trying response is returned, but that is not shown in the call-flow diagram in *Figure 1*.) The voice gateway sends out a signal initiating the call to the PSTN and informs the proxy (which in turn informs the user agent) that it is ringing the destination. When the call is answered, the voice gateway sends a 200 OK back to the proxy. This message contains the gateway's information that the user agent needs in order to set up a media stream back to the gateway. When the user agent receives the 200 OK, it responds directly to the gateway with an ACK. At that point, the call is in session, and the two endpoints pass media directly between each other. When either side wants to terminate the call, a BYE is sent from one endpoint to the other, and an ACK is returned.

In this example, once the proxy facilitated the call between the user agent and voice gateway, it dropped out of the call flow, and the two endpoints continued the session management directly between themselves. In some cases, it may be desirable for all session management to flow through the proxy. This may be for security reasons (i.e., the voice gateway is configured to accept messages only from the proxy) or because the proxy needs to be aware of all messages in order to supply certain functionality. SIP defines a "record-route" tag that is inserted in the SIP message by the proxy for such situations. All subsequent messages must then flow through the entity (or entities) addressed in the "record-route" tag(s).

3.2. Registration

SIP defines a method for a user agent to inform a server of its location. The user agent "registers" with the registrar of

FIGURE 1

Simple SIP Call Flow

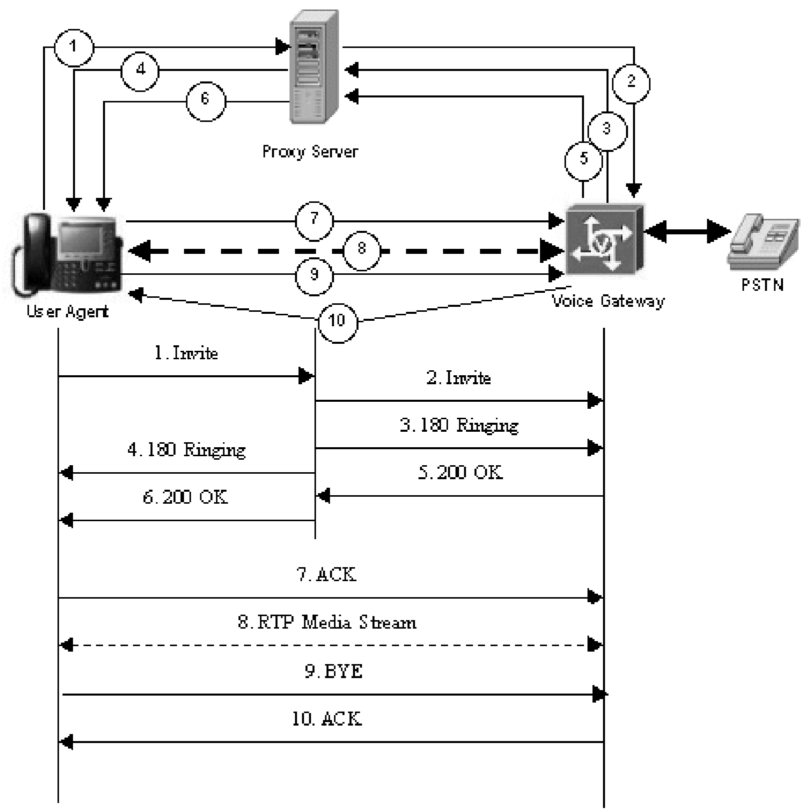

its domain, so that when a SIP message headed for that user reaches the domain, the registrar knows where to forward the message. The registration message contains a "TO" field, which defines the SIP address for which messages are intended, and a "CONTACT" field, which holds the forwarding address where the user can currently be reached. The registration is valid for an amount of time defined in the EXPIRES parameter.

The following scenario (see *Figure 2*) shows an example of a user agent registering with a registrar, and then being contacted by another user agent via that registrar.

The first user agent (*baruch@deltathree.com*) registers with a registrar and sets its contact information to baruch@123.1.2.3. Another user sends an INVITE over the Internet to *baruch@deltathree.com*, and that message is routed to the SIP proxy/registrar for the domain deltathree.com. The proxy/registrar looks up the user *baruch* in its registration database and forwards on the INVITE message to the user agent at IP address 123.1.2.3. From then on, the call flow continues as in the previous section.

3.3. Authentication

To deny service to any request that does not come from an authorized user, SIP defines an authentication protocol that allows users to identify themselves with a username and password in a secure manner. SIP supports both basic and digest authentication; basic sends the username and password in clear text while digest passes the password in encrypted form. When a user agent attempts to register or sends an INVITE to a proxy, the proxy/registrar responds with a 401 (or 407) response, indicating that authorization is required (see *Figure 3*). (If digest authentication is used, that response also contains the information necessary to build a secure transmission, which includes the encrypted password.) When the user agent receives the response, it will reissue the REGISTER or INVITE, but this time it will include the necessary authentication information (incorporating the security information just received, in the case of digest authentication). The proxy/registrar will test the username and password against its own data, and if they are valid, it will process the request accordingly.

3.4. Additional SIP Capabilities

SIP was designed to be open and flexible, and as a result it can be used for many diverse applications. One can build on SIP's location and session negotiation capabilities and apply them to various situations. In some cases, extensions have been defined that stretch SIP's functionality even further. Also, SIP is still very much a work in progress, and various additions and subtle add-ons are being suggested and worked on.

Some additional applications where SIP is being used include the following:

FIGURE 2

Registration and Inbound Call

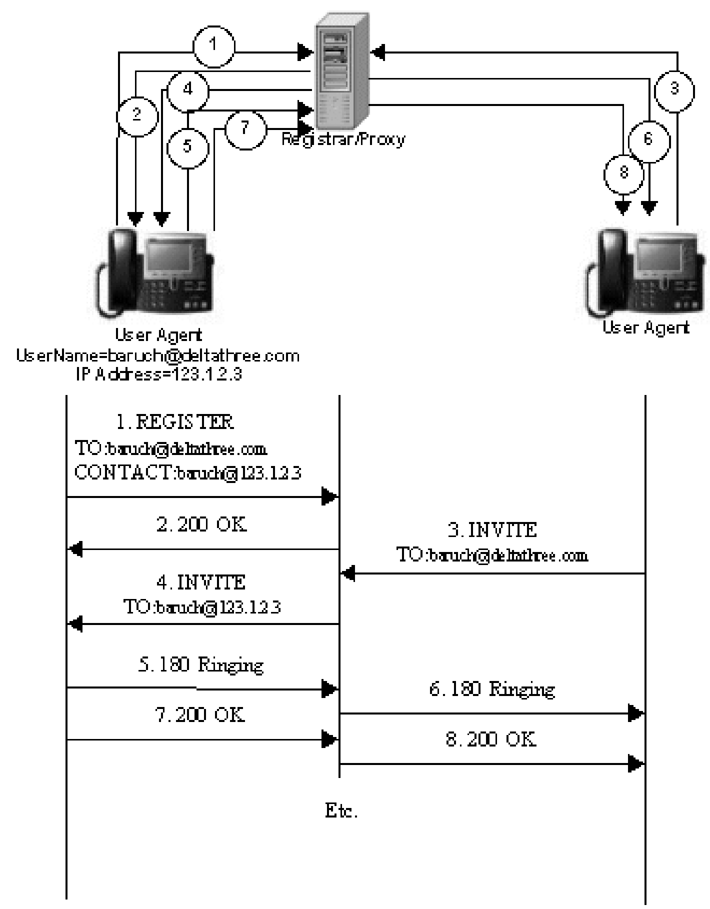

FIGURE 3

Authentication Flow

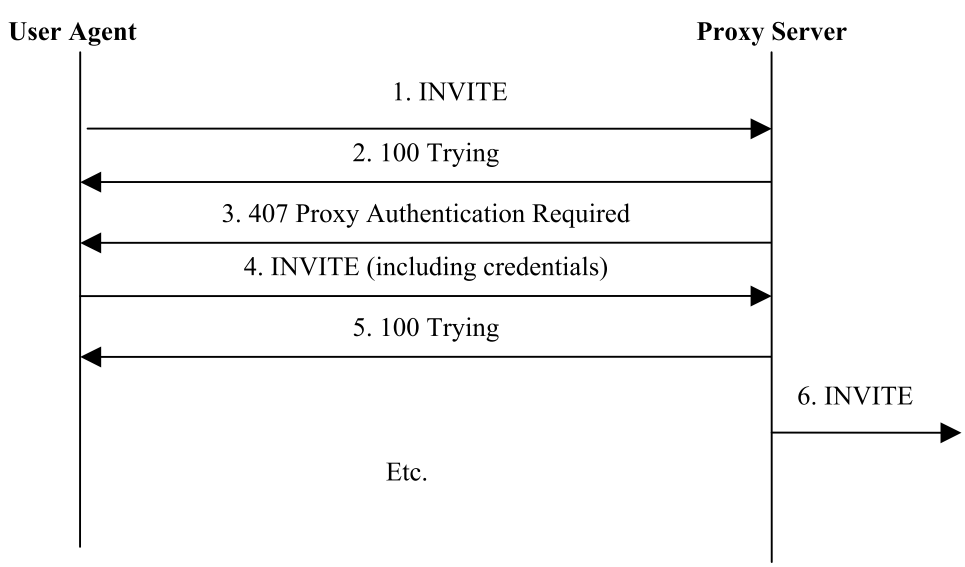

- *Presence:* SIP's registration infrastructure provides the foundation for centralized tracking of users' location and status. SIP can be used to build buddy lists and other presence-based programs.
- *Instant Messaging:* Instant messaging applications typically send short text messages between parties. SIP can accomplish this easily, since it already has a location mechanism and supports text attachments to the SIP message.
- *Event Notification:* SIP defines a Subscribe/Notify message header that is used by a user agent to request that a server or user agent inform it when certain events occur. One example of this might be an IP phone subscribing to a voice-mail server and requesting that it be notified when a message comes in. The IP phone would receive a message in the event that mail is received, and the phone might then switch on a "message waiting" indicator to let the user know that the voice mailbox should be checked.

There are numerous other examples of SIP–based applications. Any situation that requires communication over the Internet between two or more endpoints might be implemented using SIP. As SIP matures and is more widely deployed, and as developers become more aware of its capabilities, we can expect to see SIP–based applications becoming ubiquitous.

4. Pre-Paid Billing

In February of 2001, I attended a SIP conference where a well-respected expert was discussing SIP–based Internet telephony. I stood up and asked him how SIP intended to deal with billing, specifically with per minute charges. Since SIP deals mainly with sessions, per-session billing is not a problem. But, as an Internet telephony service provider (ITSP) that bridges calls between IP–based endpoints and the PSTN, our costs are based on the PSTN billing model (per minute), and so we need to charge our customers accordingly. The speaker explained that within a few years, the legacy telephony system would be completely replaced by VoIP and that, just as no one thinks of billing per e-mail, so too telephone calls would be bundled in a flat monthly service charge for IP connectivity.

Needless to say, I was not convinced, and that night proceeded to write an Internet draft that proposed a method of interfacing SIP calls with RADIUS (remote authentication dial-in user service, a standard protocol used by many billing systems). A few months later, Microsoft came out with their requirements for ITSPs wishing to terminate SIP–based traffic from their Messenger client, and per-minute charges were an essential component. As long as telephony costs are linked to usage, Internet telephony must reflect that in its billing model and must bill in real time. If not, the potential fraud risk will be enormous, as clever hackers will find ways to take advantage of loopholes, to exploit the reliance on statistical methods (such as building a rate plan based on assumed average usage), and to use any delay in updating or processing billing information to overuse an account's resources.

4.1. Real-time Billing Considerations

The majority of retail sales for VoIP applications come from personal computer (PC)–to-phone (a software client running on a PC) or phone-to-phone (via interactive voice response [IVR] gateways). Users typically sign up for these services via the Web or purchase calling cards with account information already provisioned. In either case, the model is normally a pre-paid one, where a defined amount is credited to an account and the balance is updated immediately upon termination of a call. It is essential that debiting the account balance occur as close to real time as possible, since an account must have the updated balance information available for the follow-on call. Before each call, a user's account balance is queried to make sure that there is enough credit to authorize the call. The balance is translated into duration based on the destination/origination rates for that particular user and call, and a timer is set to tear down the call if the user goes beyond the allotted duration available.

The algorithm for processing this type of billing structure is often broken into three steps: authentication, authorization, and accounting (AAA).

4.1.1. Authentication

The first stage involved in processing a call request is to retrieve an account identifier (e.g., a username) and password. In IVR systems, these are digits entered from the telephone keypad. In the case of IP–based clients, they must be passed to the server in a secure manner (in SIP this is accomplished through basic or digest authentication). Authentication means verifying that the account is valid and open for making calls and that the password matches one stored and associated with that account.

4.1.2. Authorization

The second stage verifies that the user is authorized to make this specific call (after the destination, or B-number, for the call is provided). Given a user's rate plan, account balance, and the explicit rates based on the origination and destination associated with this particular call, the billing server will return a maximum duration available. With this duration value available, the system (trusted user agent or network server) will start a timer and terminate the call if the call goes beyond the allotted duration. In the case of IVR–based calls, the voice gateway will typically have the timer functionality and the ability to close a call. In the case of IP–based clients, call teardown poses a more complicated problem, since the client usually cannot be trusted to terminate a call when funds run out.

It is imperative that the data available to the billing server for this calculation be current. If there is a delay between the end of one call and the account balance update, a user placing a follow-on call immediately after a previous one will have more credit available than should be the case. The maximum duration set for the follow-on call will allow the user's balance to run below zero, creating a liability in terms of collecting the funds owed.

I would like to stress the critical nature of fraud associated with IP–based calling products. On the one hand, the automated, on-line (and somewhat anonymous) nature of the provisioning process makes these applications particularly vulnerable to hackers and credit-card thieves. On the other hand, the costs incurred by the ITSP are real in that every minute originating from the ITSP's network and terminating on the PSTN will be charged to the ITSP.

So there is an immediate liability in terms of real money as soon as an account is opened, and a system that allows any loophole whereby a user can spend money without proper coverage will be exploited to the benefit of rogue IP telephony hackers.

4.1.3. Accounting

The last stages involved in the billing process are the recording of the call detail (CDRs), the final rating of the call, and updating the account balance. Again, this accounting must be done quickly in order for the follow-on balance to be accurate. To accurately determine the duration of the call, either the call start and stop information must be recorded and then correlated, or else at least one element in the network that can be trusted to know precisely the start and stop of a call must keep state and track the duration.

4.2. Centralized Authentication: Digest Authentication to RADIUS Mapping

SIP has no preferred method for implementing AAA functionality or billing. Indeed, since SIP is meant to manage the call session, set up, and tear down the call, it could be argued that billing is completely orthogonal to the scope and functionality of SIP. However, even in the earliest drafts, the authors of SIP began to address certain issues relating to AAA—namely authentication. Sessions originating on the Internet or other IP networks do not have the same trusted stature of a call coming from a dedicated twisted pair of copper wires that are physically attached to a device located at a definite address. In fact, the idea that authentication of the originating user is essential, and the notion that only trusted users will have access to certain services and applications available through SIP, implies either that those services may involve confidential information or that those services are billable. (Note that generally, e-mail does not require authentication of the sender, because the outbound-only nature of e-mail means that the sender is not accessing private information—only the receiving party has that constraint and so must provide a password. Also, e-mail is virtually never billed for on a per–e-mail basis.)

Whereas authentication was always considered part of SIP, authorization and accounting were not addressed. Most commercially available proxies have logging functionality that can usually be adapted for the purpose of CDR creation. Reading logs, however, implies a "batch" mindset where records are processed off-line at a later time. They are not especially suited for real-time billing applications. The fundamental requirements for real-time billing as described are a central billing service where each of the three As read from and write to the same data repository, and that information is updated transactively, that is, a follow-on call is available only once the previous call's information is processed and the balance updated. Systems working in this way need to "push" the information to the billing server at the termination of a call and not rely on writing the call information to some file or queue and waiting for the server to "pull" that information and process it.

Many legacy systems implementing real-time AAA use a protocol called RADIUS [3][4]. RADIUS is meant to provide the following functionality:

- Passing a user's password securely, not in clear text.
- Verifying that the RADIUS client is a trusted entity.
- Verifying that the RADIUS server is a trusted entity.
- Passing additional information between the client and server in a secure manner, ensuring that no one in between can tamper with the data.

To accomplish this, RADIUS makes use of a shared secret between the client and server. That shared secret must be specified to each party in a prearranged manner that is outside the scope of the RADIUS protocol.

Note that confidentiality is assured only to the password field; all other communication is guaranteed reliability (i.e., no one can change the data), but theoretically someone sniffing the data going back and forth could read the information.

RADIUS defines (among some other ancillary things) two messages: the authentication message and the accounting message. A client sends an authentication request including the username and password to the server, which responds with an accept or reject message. The accounting message has two subtypes: accounting start, sent at the beginning of a call, and accounting stop, sent on termination of a call, which includes all the information needed to create a CDR. All of these messages may also include vendor-specific attributes (VSAs), and, in that way, the client and server can pass information between them in a reliable manner. The RADIUS server will pass the information over to a billing service that reads and writes to a database and processes the information for the RADIUS server in order to determine how the server should respond to requests and what information it should pass back to the client.

The second A—authorization—requires the destination or B-number of the called party and, based on that information, determines the duration available to the user for the call. This can be implemented through the use of VSAs in either the authentication message or in the accounting start message.

As previously mentioned, in SIP, the username and password are verified using digest authentication (when security is a necessity). However, due to the different encryption and hashing (one way encryption) mechanisms used by digest and by RADIUS, there is no direct way of using the digest information and translating it to RADIUS. (The details of this are beyond the scope of this paper and can be found in [5].) We have proposed a mapping between SIP digest authentication and RADIUS that requires no changes to the SIP client but necessitates modification of the SIP proxy and RADIUS server. (These modifications have already been implemented in commercially available SIP proxies, such as Cisco Systems' CSPS product.) In this manner, a SIP proxy can authenticate a user from the data provided through RADIUS to a centralized server. Other SIP network elements may take care of authorization and accounting, again via RADIUS, to that same centralized server. Such a set-up will take care of the fundamental requirements for a real-time interface between a SIP network and a billing service.

4.3. Call Teardown and CDR Creation

Now that a real-time interface between the SIP network and billing service has been defined, the only thing still required in order to realize a pre-paid model is call teardown when the call reaches the maximum duration allotted. There are two suggested approaches as to where the element that is responsible for call teardown should be located. Each option has advantages and disadvantages, and the ultimate decision will depend on the specific network topography and other considerations.

4.3.1. Control at the Edge

The first model places the responsibility for tearing down a call at the edge—in the voice gateway that either terminates the call to the PSTN or originates the call from the PSTN. Many gateways (such as Cisco's voice gateways) have the functionality already in place to communicate via RADIUS and set timers and to tear down a call when the maximum duration is reached. The reason that this functionality has already been implemented in the gateways is that they are often used as IVR gateways for pre-paid calling-card (phone-to-phone) applications. Very minor modifications are necessary to get the same functionality to work from a SIP–based VoIP call, since all the necessary information is already present in the SIP message, including the B-number or destination, and the user information. When the call terminates, the accounting stop message is issued to the RADIUS server by the gateway passing all the CDR information back to the centralized billing service in real time.

An advantage of keeping this functionality at the edge is that the gateway is one network element that must always, in any case, keep state for the entire length of the call. It will, therefore, always be able to know the final duration of the call and can pass that information back to the billing service for processing. Without the duration, the billing process would need to correlate accounting start and accounting stop messages, adding to the processing time and resources required by the billing service.

A second advantage is that this solution scales with the gateways: As the number of ports required within a network grows, the number of gateways increases proportionally and the increased processing power needed for keeping state of all those calls is distributed over the gateways.

The major disadvantage of this scheme is that it requires that gateways be trusted elements not only within the SIP network, but also as billing entities. An ITSP may not be able to live with such a constraint if it wants to leverage least-cost routing to various partner termination (or origination) networks. If the gateway does not belong to the ITSP, then it may not be compatible with the RADIUS protocol implemented by the ITSP, or it may not be trusted to pass secure information across network boundaries. Along the same lines, there can be no billing for calls where both endpoints are untrusted entities or do not speak RADIUS. Calls from IP devices or software clients to each other—or to services such as conference servers, voice-mail programs, etc. that do not have built in RADIUS support and in any event would hardly be considered trusted entities—would not be able to be billed.

4.3.2. Centralized Call Control

The second model places the responsibility for call teardown in a centralized network element: the call controller. All SIP messages must pass through this element, which is responsible for authorization of the call (via RADIUS), for setting a timer that fires an event when the maximum duration is reached, and for sending out the SIP BYE messages to both parties when that event fires in order to tear down the call. The call controller would keep state of the call (from INVITE and 200 OK through to the BYE or timer firing) and write the accounting stop to the RADIUS server, including the call duration.

The advantage of this scheme is that any SIP call passing through the call controller can be billed regardless of the termination or origination characteristics of the endpoint. The disadvantage is that as a centralized network element, it is a single point of failure for network traffic and must be redundant, highly available, and scaleable in order to process and keep state on every call running through the network.

5. Conclusion

SIP stands poised as the method of choice for management of media transport sessions between Internet-based endpoints. Its advantages are numerous: It is highly extensible and flexible, easy to program and troubleshoot, and particularly well suited for the Internet. As such, important bodies and companies have already adopted it as their communications protocol; device manufacturers are building it into their products; and software developers are busy implementing applications and services that interoperate using native SIP.

The challenge for the service provider is to integrate SIP devices and software clients with the services and applications available and to bill for them. SIP by itself does not adequately address the problems and requirements associated with billing, specifically real-time billing in a pre-paid model that fits the particular demands and constraints of Internet users. An interface between SIP and legacy billing services is required. RADIUS offers a good prospect as the protocol for that interface, since many billing services already support it and the required mapping between SIP messages and RADIUS can be defined in a relatively straightforward manner. Development of specific billing entities that sit within the SIP network, or else some modification of existing network elements, may be required to achieve full integration between SIP devices and applications, and a real-time, pre-paid billing model.

6. References

[1] M. Handley et al., "SIP: Session Initiation Protocol," RFC 2543 (March 1999).

[2] H. Schulzrinne and J. Rosenberg, "The Session Initiation Protocol: Internet-Centric Signaling," *IEEE Communications* (October 2000).

[3] C. Rigney, "Remote Authentication Dial In User Service (RADIUS)," RFC 2865 (June 2000).

[4] C. Rigney, "RADIUS Accounting," RFC 2139 (June 2000).

[5] B. Sterman, "Digest Authentication in SIP Using RADIUS," IETF draft (February 2001), draft-sterman-sip-radius-00.

Network Outages and Quality

M. Whiting Thayer

Senior Engineer

Federal Communications Commission

Editor's Note: The opinions expressed within this paper are those of the author and are not necessarily those of the FCC.

Introduction

Virtually everything in life is governed by standards of some kind, and most people would agree that standards are a necessary part of working and living. However, standards are rendered worthless if they are not implemented and followed. And that is exactly what is happening in the telecom industry today.

For example, there are rigorous standards in place for signaling system 7 (SS7) that are not being followed. The SS7 signaling standard—established by the New York–based American National Standards Institute (ANSI), the primary body for handling telecommunications standards—was based on many months of study by numerous review boards composed of industry experts. These industry representatives found that, to ensure reliable service, diversity of many components of SS7 are required. SS7 has various reliability features for the same reason a 737 airplane has two engines—for backup. If an airplane had only one engine (and it can be airborne with only one), most passengers probably would be quite resistant to flying in it.

The primary SS7 standard relating to diversity—these standards are readily available either directly from ANSI or from the Alliance for Telecommunications Industry Solutions (ATIS) document center (www.atis.org)—is ANSI T1.111-1996. This standard was reaffirmed last year and will be ANSI T1.111-2000 when it is published. However, there are companies implementing SS7 without putting the required diversity in place, thereby rendering the standard meaningless.

Arguments Against Standards Compliance

A common argument against compliance with the established standards is that they are "voluntary." And this is true. In fact, the voluntary nature of standards is common over all industries, not just telecommunications. The concern, then, becomes whether these companies are selling products or services, such as toll telephone service based on a standard, even though they are not following the reliability aspects of the standard. That is misrepresentation. Moreover, such violations are not immediately visible. No one can tell just by picking up a telephone that a certain standard has or has not been followed. It is only when a failure occurs and customers are not able to make or receive toll calls that light is shed.

Some also argue that these standards may not be cost-effective. Many carriers are making a conscious business decision not to implement the standard requirements for diversity, using this reason (of not being cost-effective) to justify their decision. Considering some of the problems carriers are having, cost-effectiveness depends on what costs are included. Some costs, after all, are difficult if not impossible to quantify. For example, what is the cost to the public of not having the telephone service they are paying for? What is the cost to people that are unable to dial 911 and reach an emergency center?

Best Practices

The industry also has what are called "best practices." Following some severe SS7 signaling outages around 1990–1991, the Federal Communications Commission (FCC) asked the industry to meet, examine existing outages, and develop best practices that would, if followed, eliminate common problems. They have done and still do this. Best practices are not official standards, in that they do not have to go through the rigorous review process required for ANSI standards. No rules state that best practices must be followed. Best practices represent the industry's own consensus as to what would help make a reliable network.

One of the industry organizations is the Network Reliability Interoperability Council (NRIC), a two-year federal advisory committee. As an example of their work, Focus Group 3/Subcommittee 1 of NRIC–4 examined all of the existing best practices that had been established over the years, weeded out the ones that were obsolete because of technological changes, and developed new best practices where there were gaps. This was done in response to some network outages where obvious failures had taken place.

Another group, ATIS, was formed right at divestiture by the local telephone companies. It has since expanded, taking in a host of other groups, including manufactures, interexchange carriers, and others. It sponsors the Network Interconnection Interoperability Forum (NIIF), which has developed a reference document ("Installation, Testing, and Maintenance Responsibilities for SS7 Links and Trunks—Attachment G: Link Diversity Validation Guidelines," which is also available from the ATIS document center) that covers many aspects of network reliability that should be managed. It basically outlines a set of best practices to achieve the SS7 diversity requirements spelled out in the ANSI T1-111 standard. This document was put together by the industry, not by the FCC.

Outage Reports

Per Section 63.100 of FCC rules, any carrier that experiences an outage of at least 30,000 lines for 30 minutes or more must report that outage to the FCC. These are just telecommunications outages, not Internet or packet switching outages. The newest NRIC, NRIC–5, is working on what reports should be required for packet-switching networks. Many best practices are identical between regular telephone and packet networks.

Organizations that experience an outage must file a preliminary report within hours of the occurrence and a final report 30 days later. The preliminary report need only say something to this effect: "We have a problem. We do not know why. We do not know how long it will last." While this initial report may not seem to be particularly useful, it alerts the FCC if someone from Congress were to call and ask if they knew there was an outage in such and such a place. Reports on all the outages that have occurred since 1996 are posted on the Web at www.fcc.gov/oet/outage and available to the public. There is no summary or analysis available there, but there is a scan of the actual paper reports that were filed by the carriers.

There are approximately 200 outage reports filed every year, which averages out to be nearly four per week. These are not isolated instances, and there is great concern that many outages are not being reported. For example, if there were 29,000 lines out for a week, or 10,000 lines out for a month, they would not be reported because they are below the 30,000 line/30-minute requirement. The few exceptions to this rule are when outages occur from fires that affect more than 1,000 lines, outages at major airports, and outages affecting nuclear power plants.

Common Causes of Outages

A common industry scapegoat reason for outages is backhoes—road workers digging up the street and cutting the cables. In reality, though, the majority of outages are caused by problems inside the carrier's building, not out in the street. In other words, the majority of outages are caused by human errors—employees of the reporting company or contractors putting in new equipment or modifying existing equipment. These kinds of outages are called procedural outages, and they have been increasing in frequency.

Many of the outages are directly related to the lack of diversity in the network. The SS7 standards and best practices call for diversity in almost all aspects of SS7—digital timing supplies, power supplies, digital cross-connect machines, cables, cable routes, etc. If that diversity is not there, the standard suggests that the maximum annual downtime requirements will not be met. A lack of diversity does not cause outages, but it does significantly contribute to the extent of the outage (number of customers affected). In other words, when something fails and there is little to no diversity, there is a far greater chance that a much wider area will be affected by the outage than if there had been diversity. Again, the single-engine 737 analogy works here: the fact that it only has one engine will not cause a problem in and of itself, but this will make the plane more vulnerable to other problems.

Many outages are single-point failures. This means that one point in the network has a problem—one cable gets cut, one switching machine goes down, one network component fails, etc.—and the resulting problem affects a very wide area.

Another common failure is the simplex-duplex failure. In other words, many telephone network components are built with duplexed components, so that if one fails, it automatically switches to a backup, and service is not affected. Synchronous optical network (SONET) rings provide a good example of this type of failure. If a SONET ring has a failure on it, it automatically switches. No customers are affected by a SONET ring going into a simplex mode; in fact, that is the whole point of this kind of architecture. However, if the simplex problem is not repaired in a timely fashion and there is another failure, whatever or whoever was served by that component will be out of service. In one case, two and a half months elapsed between the time a SONET ring went into a simplex mode and when another failure occurred elsewhere on that same SONET ring. Since the first failure had not been repaired, all services provided over the SONET ring were affected.

Network Architecture

Figure 1 offers an illustration of telephone-company buildings (represented as the dots) and the connections between those buildings (the solid lines). This is a physical network, not a logical network. Logically, there could be a connection between the two upper dots. Physically, the traffic route goes down and through two other buildings.

This kind of configuration does not apply to rural areas only. The diagram in *Figure 1*, in fact, is a simplistic version of the Second Avenue fire that occurred in New York City back in the 1970s. The dot on the left-hand side of the line connecting the two halves of the drawing could represent the Second Avenue Building in Manhattan, through which most traffic between Manhattan and Brooklyn traveled. The building burned up. And though it did not affect the rest of the network, little traffic went between Manhattan and Brooklyn until other paths could be established.

It also must be remembered that an entire building does not have to be destroyed for an outage to occur; it could be just one common piece of equipment. And with this old architecture, if one dot (building) goes out, the rest of the network is still fine: calls can still be completed between customers served by any of the buildings that were not directly impacted.

Figure 2, representing SS7 architecture, is another simplistic diagram. This shows only one end office, toward the bottom of the diagram. The dotted lines represent packet-switching circuits. They go up through one of two mated signal transfer points (STP), and these are all interconnected. There is total diversity between these paths—or there is supposed to be—and the STPs are connected with diverse routes to two service control points (SCP). The SCP

FIGURE 1

Architecture (Old)

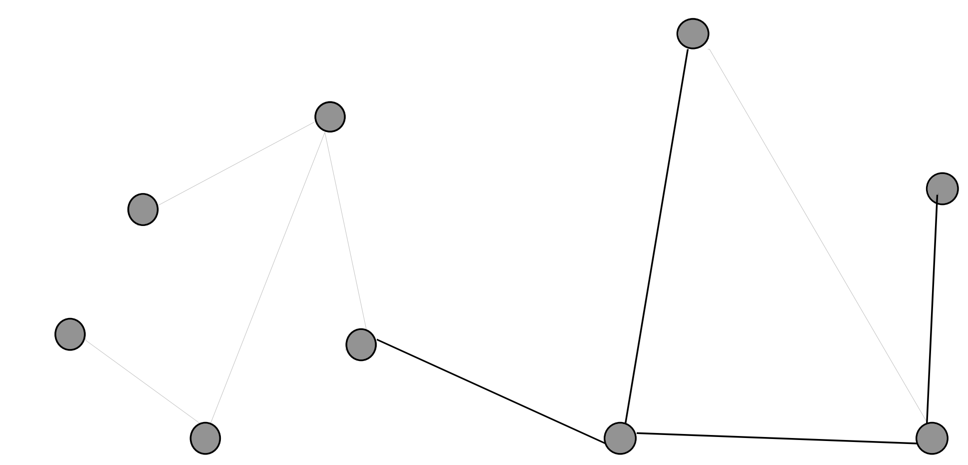

is sometimes called "The Godfather." No traffic moves unless The Godfather says it can. If no one can talk to The Godfather, traffic will not move between buildings. This is how SS7 should work. Furthermore, none of these components should have more than 40-percent utilization. This way, if anything fails, the mated component can carry the entire load without a problem.

Today's architecture is represented in *Figure 3*. The building (dot) in the diagram on the left with the arrow would have two diverse paths (dotted lines), assuming the two STPs are on the right-hand side. Obviously, the administrators of the network provide as much diversity as possible; but there are two buildings that are common, as well as the path between them. In fact, the signaling paths from all five buildings on the left side of the diagram would have to go through the same two buildings and the common path between them to reach the STPs. Thus, there exists several places where there is a single-point failure potential that could restrict the customers in all five buildings on the left from the ability to make or receive calls from any other building in the network. In other words, with today's architecture, the signaling control messages for all inter-building traffic must be able to reach the SCP (The Godfather) before traffic can flow.

Notes on Common Channel Signaling

Some valuable points for better understanding common channel signaling (CCS) can be gleaned from *Bellcore Notes in the Networks*. A few of the most noteworthy quotations from this report are included in the following sections.

On CCS Architecture
"A crucial part of designing an SS7 network is including sufficient equipment redundancy and physical-route diversity so that stringent availability objectives are met" (section 6, page 255).

On CCS Availability Objectives
"Each network access segment should be down (an average of) no more than 2 minutes per year" (section 4, page 50).

"The backbone network segment should be down a negligible amount of time (that is, close to 0 minutes downtime per year)" (section 4, page 50).

On CCS Diversity
"The above allocation assumes an ANSI–based reference architecture with two-way diversity for the A-link sets and three-way diversity for the B-/D-link sets." (section 4, page 50)

FIGURE 2

Architecture (SS7)

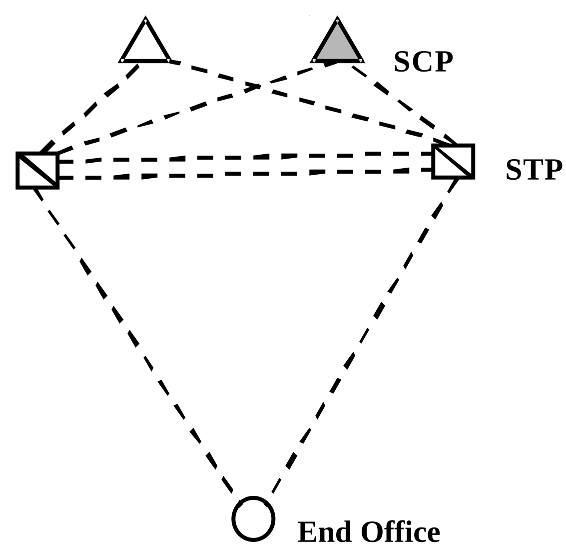

FIGURE 3

Architecture (Today)

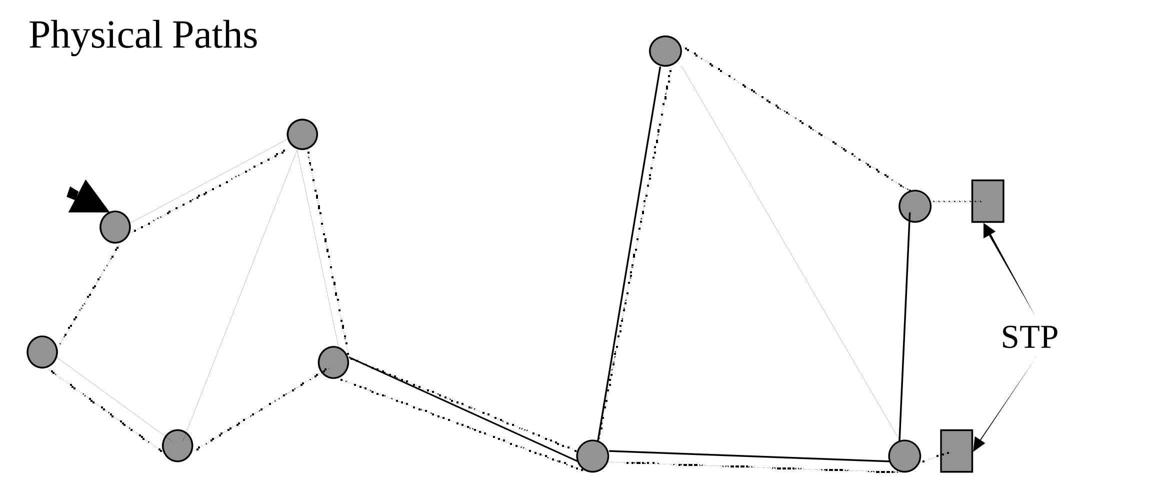

Problems with Diversity

Diversity problems are twofold. First, there may be only a single physical path available for part or all of a route. Because it is only a single physical path, everything goes over that path just like on the previous diagrams. There is no way to provide diversity even though the standard calls for it until such time as a diverse physical path is established.

The second problem with diversity is administrative. The tools that the carriers use to assign circuits to the various network components were not designed to handle diversity, making it difficult for the people who assign circuits to place and maintain them in a diverse fashion. In many cases, circuits are moved from one path to another without recognition of the diversity implications.

Case Studies

The following section offers a few illustrations from the 2000 fourth-quarter outages. The full reports are available on the FCC Web site at http://www.fcc.gov/oet/outage.

- ***Outage 00-168: Houston, Texas:*** More than 718,000 customers in 17 wire centers [that is, telephone-company buildings] suffered more than 2.6 million blocked calls. Both E-911 and SS7 links were knocked out of service for two hours and twelve minutes. It was a hardware failure in one building. It was also a simplex-duplex failure. So when the initial trouble was not repaired, the duplex side failed, and everything was lost.
- ***Outage 00-191: Silver Spring, Maryland:*** On November 13, almost two million customers were cut off for almost four hours due to digital cross-connect bay failure. It knocked out almost 200 DS–3 or T3 circuits. Everything went through one piece of equipment in one building and had a very wide impact. After a momentary power failure, a surge blew the fuses, which were the wrong size.
- ***Outage 00-192: Nashville, Tennessee:*** The following is a good example of the STP diversity. The power supply to the STP failed, so the STP went down. This would not have disrupted service if the links were loaded to 40 percent of their capacity, but somebody had grossly overloaded the links. It was a simplex-duplex failure because the initial overload was not fixed. When Nashville failed, the STP in Atlanta could not handle the load, resulting in nearly 185,000 blocked calls for over seven hours.
- ***Outage 00-194: Montgomery, Alabama:*** A tandem switch serving eight wire centers plus independents failed. This made 34 public-service answering points (PSAPs) not available to answer E-911 calls for 42 minutes. This was the result of another fuse problem: the power-transfer switch fuses were the wrong size.
- ***Outage 00-219: Miami, Florida:*** Cellular traffic is not immune to such outages either. Cellular traffic goes from the tower to a cellular switch, which is connected to a traditional network tandem switch. The links went down here between the cellular switch and the tandem switch for almost three hours. A SONET multiplexer went bad.
- ***Outage 00-220: Klamath Falls, Oregon:*** Thirty-one thousand customers ended up losing E-911 service because there was no diversity for the E-911 circuits—something that violates best practices. However, there was diversity in the SS7 circuits. Thus, customers did have long-distance service for making and receiving calls, but they could not call 911 for more than 16 hours. The cause: vandals had cut a fiber cable.

Introducing Software Paradigms to Hardware Design

Kjell Torkelsson
Vice President, Engineering
Celoxica Ltd.

Background

With today's complex systems and high-performance silicon, the worlds of hardware and software engineering are converging fast. Traditional processor-based solutions start with a high-level design, which is partitioned into hard and soft components.

For hardware engineers, taking complex functions and describing them in hardware description language (HDL) to produce a high-performance field programmable gate array (FPGA) or application-specific integrated circuit (ASIC) design can be extremely challenging and time consuming.

For software engineers, these same functions are relatively simple to implement in code that can easily be integrated with the overall system. However, the sequential nature of processors and their reliance on high clock speeds to perform fast calculations often imposes limits on performance.

The ideal would be a middle ground, with the performance advantages of parallelism gained though hardware and the ease of design and integration of software. A few years ago, this would have been a pipe dream, but advances in software tools and the growing functionality and capability of field-programmable devices now enable software concepts to be implemented on hardware, without getting involved in a time-consuming and costly HDL design loop.

By using modern, multimillion gate FPGA technology, software components can be realized in silicon faster than traditional methods, enhancing the time-to-market opportunities. What is more, such solutions can extend the lifetime of hardware, as new functions or efficiencies can simply be reprogrammed on devices out in the field. Programmable gate arrays can even be modified over the Internet.

Bridging the Gap from Software to Silicon

Celoxica's recently launched DK1 design suite compiles from Handel-C to an FPGA net list without requiring an intermediate HDL step. It allows system developers and software engineers to accelerate algorithms by expressing them in parallel and mapping them directly in hardware. The benefits are especially acute in areas such as broadband networking, third-generation (3G) base stations, or Internet backbone routers, where complex software algorithms such as encryption or compression need to be run at higher and higher speeds.

For example, in networks or high-traffic Web sites, the slow throughput available with software-based encryption is a major performance bottleneck in completing secure on-line transactions. Implementing the triple data encryption standard (DES) algorithm in FPGA-based hardware using the DK1 design suite provided throughput in excess of 6 gigabits per second (Gbps). What is more, a single software engineer was able to complete this project from specification to hardware in only one week.

Sequential algorithms are frequently best performed by software implementations using conventional central processing unit (CPU) architectures. Optimal hardware acceleration comes most often from functions performed in parallel, typically in the physical form of a hardware co-processor closely coupled with the CPU.

Celoxica's specially developed programming language, Handel-C, takes this into account, providing the extensions required for hardware development. It is firmly grounded in ISO/ANSI-C but with an additional statement for parallel processing (par) as well as constructs for communication between parallel elements. The language also provides flexible data path widths, versatile memory architectures, and interfaces to external hardware.

Other C-based design methods have emerged to support the need for high-productivity, system-level design at the algorithmic level. These fall into two camps: behavioral methodologies and RTL–based methodologies.

Behavioral compilers give the user little control, automating the processes between input and output in a way that has been likened to a "long, dark tunnel" that has a tendency to generate code that is too big and too slow.

C-based RTL methodologies tend to provide more control but wrap VHDL and Verilog functions in clumsy versions of

C or C++ and hence do little to remove the low-level complexity of RTL design.

In contrast, Handel-C delivers a level of abstraction over RTL but maintains manual control for more efficient code than behavioral approaches allow. Design engineers need both sequential and parallel expression of functions from their tools. RTL, which expresses parallelism by default, require low-level coding of state machines to implement sequential expressions-a laborious and verbose process. Handel-C combines these two worlds more elegantly with a sequential language, able to express parallelism by using simple par statements to add parallelism where it is needed.

Handel-C and the DK1 development environment allow designers to use software paradigms and library code and even to import legacy HDL blocks. In exactly the same way as newly written code, these can be simulated within DK1 and co-simulated with other tools. These facilities provide huge flexibility, speeding design by code reuse, and make a modular approach to design practical. In addition, because it is so simple to change code and re-simulate or even update the target FPGA, engineers can work using the familiar process of iterative improvement to individual functional blocks and are free to try creative approaches to gain improved performance, without the fear of dramatically extending delivery time.

Timing and Control

Handel-C is designed around a simple timing model that makes it highly accessible to system architects and software engineers. Each assignment in the program takes one clock cycle to execute, giving the programmer full control over what is happening in the design at any point in time. Results are predictable and controllable, as easily in software as in the final hardware.

Use of the par statement is illustrated by the following program extract:

```
/*******************
* The main program *
*******************/
void main (void)
{
   par
   {
      SyncGen();
      Display();
      PerFrameUpdate();
   }
}
```

The program described might be the core of a video "bat-and-ball" game. The screen consists of a background color surrounded by a border and a moving square (the "ball"), which moves in straight lines and bounces off the border. At the top level, the program consists of the parallel composition of three processes.

This fragment highlights the aforementioned major difference between Handel-C and conventional sequential languages such as C-that of parallelism. This program calls for three separately executing processes to run concurrently. The first one is the SyncGen process, which is responsible for producing standard video graphic array (VGA) video sync signals and for updating the ScanX and ScanY variables, which track the current pixel position both during display and during the hidden parts of the video signal, namely the horizontal and vertical blank intervals. The other two processes synchronize themselves to this one by inspecting the values of these two variables. The Display process is responsible for producing the video output signal, and the PerFrameUpdate process changes the variables used by the display process to advance the state of the "game."

There are many other features that make Handel-C a highly relevant language for developing functionality that will be deployed in FPGA co-processors. Firstly, it supports complex C functionality, including structures, pointers, and functions (shared and inline), thus ensuring a shallow learning curve for software engineers and allowing rapid implementation of very complex, modular systems. With extended operators for bit manipulation, and high-level mathematical macros (including floating point), it allows rapid translation of digital signal processor (DSP) algorithms to efficient hardware.

With Handel-C, there are no state machines to design. Control flow comes from C statements such as if, case, and while, so complex sequential control flows can be designed simply, in a way that is intuitive to software engineers. And the simple and consistent syntax extensions for specific hardware features such as RAMs/ROMs, signals, and external pin connections provide for efficient use of available hardware without cumbersome syntax.

Handel-C programs can automatically deal with clocks, clock enables, and synchronized data transfers across clock domain boundaries. The result is a major reduction in complexity, plus the use of the optimal clock rate for each part of the design enables increased speed and reduced power consumption.

Once the algorithm has been optimized, it is simply a matter of compiling the software down to hardware. Celoxica's DK1 produces an architecture-optimized EDIF netlist, targeted on either Xilinx or Altera programmable logic layout tools. Tests have shown that the compiler is efficient both in terms of synthesis time and the area and delay of generated designs. Benchmarks against typical VHDL designs show comparable, if not superior results, on designs produced with Handel-C. *Tables 1* and *2* show a few such examples.

DK1 includes a profiling tool that gives area and delay estimates, which are presented as color-coded annotations of the source code. This makes optimizing the generated hardware extremely simple, even for relatively inexperienced engineers. In addition, users can supply timing constraints at the source level, which are automatically translated for the target technology. At the lowest level, DK1 generates readable signal names for tracing back through post-layout timing analysis.

As well as timing analysis, the debugging environment provides in-depth features normally found only in software environments. These include symbolic debugging using breakpoints, single stepping, watched variables, and the ability to follow parallel threads of execution.

TABLE 1

Benchmark 1-DES Encryption Block

Design	Language	Clock (MHz)	CLB Slices	Rate (Mbps)
Sequential DES (16-cycle latency; one 64-bit block/16 cycles)	Handel-C	105	267	420
	VHDL	101	255	404
Fully Pipelined DES (16-cycle latency; one 64-bit block/cycle)	Handel-C	101	3025	6464
	VHDL	32.5	2528	3,052

TABLE 2

Benchmark 2-A Time Domain Multiplexer for a Multimedia Gateway

Language	Clock (MHz)	CLB Slices	Logic Depth	Lines of Code	Development Time (Days)
Handel-C	39.9	694	7	1,499	24
VHDL	34.9	1,032	9	1,748	43

Co-simulation and verification facilities are built into the Celoxica tool chain, facilitating co-design with instruction set simulators, VHDL simulators such as ModelSim, and external C test benches. A key benefit is that hardware/software partitioning decisions can be made much later in the design cycle.

In the Real World

Whether for rapid prototyping or first-generation final products, this approach gives a head start in time to market, with lower risk and the potential to upgrade products in the field. So much for the theory, but what about practical experience?

One customer halved the time it took to implement a signaling system 7 (SS7) protocol, used in telecommunications systems, using Handel-C rather than VHDL, producing more compact and efficient FPGA code in 24 days, compared with 43 days. Another company, designing an IPv6 router, developed core components in just four man-months of software-engineering effort with DK1, while its team of hardware engineers worked for 12 man-months using Verilog and still failed to produce a working design-the working Handel-C code ran to 40 pages; the incomplete Verilog design had reached 200 pages. A third customer ran a competition between a traditional hardware design team and a small group of software engineers using the Celoxica technology, with the aim of creating an MP3 audio processor. In only seven weeks, the latter successfully converted existing software and produced optimized working hardware that beat design specifications. In the same time period, the hardware group using HDL had not even finished writing the specification.

Success in these and other trials is providing the confidence for software engineers to seize the opportunity and, where appropriate, extend their domain into the hardware field. With technologies such as this, there is little to lose and a great deal to be gained.

Appendix: Writing Code for Optimized Hardware

In Handel-C registers are implemented using flip-flops and all other circuitry (forming the logic making up the content of assignments) is made up of logic gates. Each of the gates in the circuit has delay associated with it as the inputs propagate through the gate to the outputs (there is also a delay associated with the interconnecting wiring on programmable logic). An average delay for a four-input gate on a modern FPGA is of the order of a few nanoseconds, which seems very quick. However, it is important to ensure that the total logic delay for each assignment does not exceed the period of the circuit clock. We have chosen 20 MHz as the rate for a boat game, resulting in a maximum propagation delay of 50 ns.

Understanding the Logic in an Assignment

Figure 1 is a simplified diagram of an FPGA with some internal circuitry, interfaced to an external RAM. Paths 1 and 2 correspond to the following two statements in Handel-C:

FIGURE 1

Simplified Diagram of an FPGA

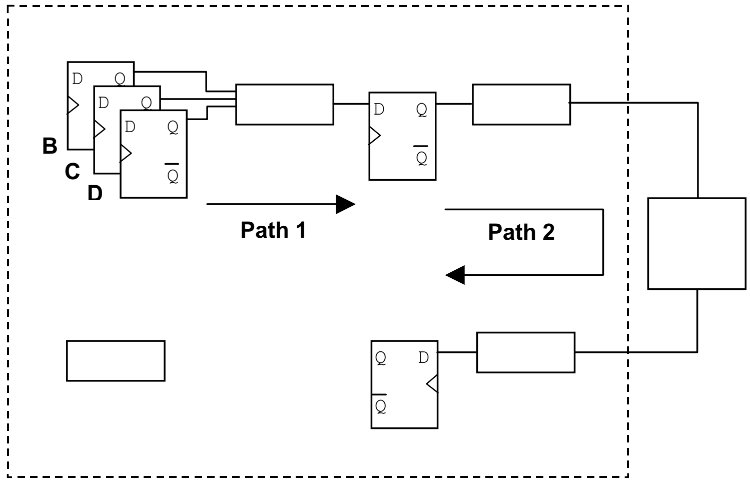


```
seq
{
A = B * C + D; /* Statement One */
E = SSRAM0[ A+1 ] - 1; /* Statement Two */
}
```

The first statement is implemented entirely in logic internal to the FPGA. This form of statement allows the most complex expressions, but the delay through the logic must never exceed 50 ns. The second statement involves an external RAM access. Variable A passes through logic to increment it by one, and the resulting value forms the address input to the external RAM.

There is then a fixed propagation delay through the external RAM from address pins to data pins (typically about 10 ns). The data re-enters the FPGA, has one subtracted from it, and is finally latched into another register. This total path must also not exceed 50 ns. To ensure that you write logic that is sufficiently fast for your chosen clock rate, it is important that each expression being assigned is kept relatively simple.

RAM Access

To minimize the logic for external and block RAM accesses, the RAM address is supplied directly from a register, and the data from the RAMs is put into a second register. If the same RAM is accessed from many places, Handel-C builds logic to multiplex the address lines for each access out to the RAM address pins. This logic has its associated delay, so the number of places at which RAM (or any shared resource) is accessed should be kept to a minimum.

Reducing Logic Depth

The code fragment

```
seq
{
A = B * C + D; /* Statement One */
E = SSRAM0[ A+1 ] - 1; /* Statement Two */
}
```

would have shorter logic paths written as

```
seq
{
X = B * C;
A = X + D + 1;
Y = SSRAM0[ A ];
E = Y - 1;
}
```

The second code fragment would then take four clock cycles rather than two. In most cases, however, we need to be able to perform one complete calculation every clock cycle. The number of clock cycles can be reduced to one by using *pipelining*.

Pipelining

The following code snippet illustrates an expression that is potentially too complicated for the application constraints.

```
a = (b + c + d) * (e + 1);
```

As before, this can be split into a number of sequential assignments, each with a logic depth shallow enough for the timing requirements. To get a single calculation per clock cycle, the statements can all be evaluated simultaneously by using *par*.

```
while(1)
{
par
{
/* First clock cycle */
sum1 = b + c + d;
sum2 = e + 1;


/* Second clock cycle */
a = sum1 * sum2;
}
}
```

After the first clock cycle:

- Correct values for sum1 and sum2 are calculated.
- The value in a depends on their previous values.

After the second clock cycle:

- The full result is available in register a (*simultaneously* the sums have been calculated for the next values of b, c, d, e).

By the next clock cycle:

- There is a new value for the calculation.

This shows that after a delay (or *latency*) of one clock cycle, a result for the calculation is available every cycle.

Pipeline Coding Requirements

Pipelined code can be slightly larger than equivalent simple code, but the logic delay should always be considerably lower. Pipelining allows you to perform very complex operations at high clock rates. However, pipelined code is marginally harder to write because there is a delay before the data is available. When a pipeline is started, the first few cycles of output must be discarded until the true data propagates through.

When writing code that uses the output from the pipeline, you must take care not to use some of the junk data or discard some of the initial valid data.

Operator Complexity

There is a hierarchy of complexity for Handel-C operations, with (for a given bit width) division being more complex than multiplication and multiplication being more complex than addition. Moreover, the < and > operators are more complex than an == comparison. Thus, divisions should be avoided unless absolutely essential. Wide multiplications should, whenever possible, be the only operation in a particular assignment. It is also best to use == rather than < or > for any Boolean operations.

Conversion to Fixed Point

Wherever possible, it is more efficient to use fixed rather than floating-point operations in hardware. In most cases, floating-point hardware is easily replaced with fixed-point, where a fixed number of the bits in the representation are to the left of the decimal point and a fixed number are to the right. You can set the number of bits to the left and right of the decimal point to different values for different stages in an algorithm to provide the accuracy required as efficiently as possible. In cases where high accuracy is required, it is important to calculate the possible error incurred by converting from floating-point to fixed.

New OSS Software Can Help Service Providers Regain Control of Their Businesses

Catharine Trebnick
Senior Vice President, Marketing
CPLANE, Inc.

The telecommunications services market has certainly experienced its share of growing (and contracting) pains over the last 12 months. The marketing promises of the late 20th century have given way to the business realities of the 21st century, and many carriers and service providers have been sacrificed along the way. It is not surprising then that those who have survived are taking a more cautious and conservative approach to managing and growing their businesses.

Nowhere is this more evident than in the carrier spending plans tracked by industry observers. In 2000, according to a recent report from Salomon Smith Barney, U.S. carriers spent more than $97 billion in capital expenditures (CAPEX), with most of it going to network infrastructure upgrade and expansion. In 2002, CAPEX spending will likely fall to around $55 billion—a $40 billion decrease. There are two primary reasons for this reduction. The obvious one is current economic conditions, but the more interesting one is the shift in business focus from operational expansion to operational efficiencies.

If you look more closely at CAPEX spending plans, you'll also see plans for increased spending on software—particularly operations support system (OSS) applications that will help to improve operational efficiencies and extract greater value from the business. These applications do a great job handling discrete processes such as order management, service assurance, and fault management.

Unfortunately, these applications do not address a fundamental control problem that service providers must solve—that is, the ability to effectively control both service and network resources. After all, optimal service performance for the end customer and optimal use of network resources for the service provider are both essential to the survival and growth of the business.

The good news is that a new class of OSS software—which combines service activation with network optimization—is emerging to address this issue. Available from several legacy OSS vendors and next-generation upstarts, the software is already generating a lot of interest from regional Bell operating companies (RBOCs) and next-generation carriers alike. Unlike many other OSS applications, these new activation/control applications provide a more holistic level of control by bridging the gap between three layers of the telecommunications management network (TMN) model: the service-management layer (SML), the network-management layer (NML), and the element-management layer (EML) (see *Figure 1*). The software's appeal is that it empowers service providers with the ability to effectively control both service and network resources in a way that is consistent with strategic business policies and objectives. And that means control of their business.

In the remainder of this article, we'll briefly look at how this software bridges these layers and provides a complete OSS activation and control solution. Let's start with the SML and the requirement to rapidly create new revenue-generating services.

In the past, the definition of new services was most often tied to changes in the underlying network infrastructure; in fact, investments in the infrastructure typically drove the creation of new services. As a result, the process involved in developing new services could be complex, time consuming, and expensive. This process also tended to stifle service innovation and prevented service providers from being able to quickly prototype and test new service offerings. The new software promises to empower carriers with the tools to respond faster to market opportunities with new services.

The software allows providers to directly develop their own service-level specification (SLS) as a high-level extensible markup language (XML) document. The SLS document typically contains information on the customer, service type, service termination points, and service connectivity attributes—things such as bandwidth, response time, latency, security, etc. Some or all of this information may be

FIGURE 1

Caption: A new class of software is enabling service providers to more effectively control the key service and network resources that drive their business models.

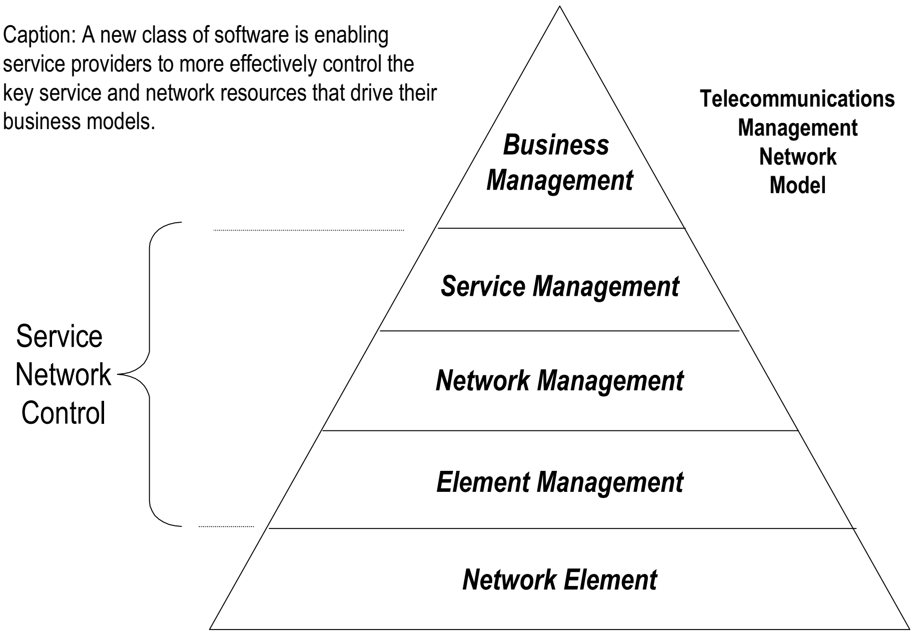

retrieved from or shared with other OSS applications handling order management or service assurance.

These SLS connectivity attributes are often key components of customer service-level agreements (SLAs), and customers pay a premium to ensure that their SLAs are met. That means that the service provider must work hard to meet those expectations. So, while operations at this layer focus on "what" the service will deliver, it is even more important that the software understand "how" the network will deliver it.

It is at the NML that an understanding of the network comes into play. In fact, this intersection of the service and the network is where the issue of control becomes business-critical, and this is one of the critical gaps in the TMN model that service network control software helps to bridge. One of the main objectives of this software is to enable service providers to better understand the impact of new services on the network, so that they are better able to determine how best to use expensive network resources to support revenue-generating services.

Therefore, at the NML, the service network control software must be able to translate the high-level service request into a set of instructions that the network can perform to ensure that the customer SLA is met. That's the real objective—not just blind activation of a service request, but optimal delivery of the service across a complex network infrastructure.

Getting control of network resources would be difficult enough if the network was composed of equipment from one vendor or one technology But many service-provider networks have been built over many years with new and legacy equipment operating at multiple network layers (optical, asynchronous transfer mode [ATM], Ethernet, Internet protocol [IP], etc.).

The software addresses this complex problem by first extracting data from other OSS tools—such as topology information from a network-management system (NMS)—to build a comprehensive information model that provides real-time data on the physical and logical topology of the network and on how resources are currently being used. This information, combined with analytical intelligence built into the software and the vendor-independent control-plane architecture of the software itself, allows the service provider to determine which technologies and resources are available to fulfill the requirements of the service.

For example, assume that a carrier was rolling out an IP virtual private network (IP–VPN) service to link multi-site enterprises between various metro networks across the United States. In that case, it may wish to perform traffic engineering within the core of the network to set up multiprotocol label switching (MPLS) label-switched paths (LSPs) between each metro network. That level of traffic engineering would enable the carrier to configure specific MPLS–enabled network elements to ensure that they are individually tuned to support the IP–VPN service. That, in turn, ensures optimal use of key resources *and* performance of the service, *and* satisfaction of the customer.

Preparing the network for optimal delivery of premium services obviously requires more than just engineering the core of the network. To begin with, any request for network resources must be evaluated against admission criteria—policies established by the service provider. Based on these policies, the software can determine whether or not a request is authorized to use the resources and, if so, whether resources are available to support it. This is a service-layer extension of well-understood "on-the-box" control functions found in protocols such as private network-to-network interface (PNNI) or resource reservation protocol (RSVP). It gives the provider the power to arbitrate access to resources directly through high-level business rules, rather than rely on arbitrary decisions made by devices that are unaware of a network-wide provider goal. If a request is admitted, it must be mapped onto appropriate resources in the network.

Implementing any service requires resource controls in the data path. For example, traffic entering the network might need to be policed to an agreed rate or shaped for injection into the network—this is known as "traffic conditioning." Deeper into the network, shared resources must be partitioned between multiple traffic types to ensure that each receives the appropriate service level (bandwidth and quality of service [QoS])—this is known as "bandwidth management." And, as discussed, in the core of the network, fat pipes (trunks, or wavelengths) must be in place to deliver sufficient capacity to the aggregate traffic load.

The output of the various operations in the NML is a "network-level specification" that is linked to the EML. In the past, a specification like this was created by hand and delivered to an army of network engineers responsible for manually configuring each network element along the data path. This was a time-consuming operation obviously prone to human error. When the configuration work was complete, the next step was often troubleshooting to identify the source problems that would no doubt occur along the way.

With this new OSS software, this process is largely automated, with network engineers involved at critical checkpoints. Instead of a human-controlled process, it becomes a software-controlled process where the software works directly with the vendor element-management systems (EMSs) to ensure that specific network elements are configured to support specific connectivity attributes of specific data services. Errors are minimized, if not eliminated, because the network specification is based on real-time data extracted from the information model. As mentioned earlier, this database tracks resource availability and service traffic flows, and then determines how best to use available network resources to support service requirements.

Once the appropriate set of network elements have received instructions on how to behave, the new service is then deployed across the network. Customers can then be activated in near real-time and feel confident that the performance of the service will be in line with their expectations and SLAs. In fact, data in the information model can be provided to other OSS applications—such as service assurance and billing—to "close the loop" with the customer and provide evidence that the SLA has been met.

The flexibility of this new breed of software also allows enterprise customers to gain direct, Web-based access to their service orders, so they can make instant changes and upgrades to their services and have those changes activated in near real time.

Now more than ever before, service providers are looking for ways to achieve the "faster, better, cheaper" objectives of any successful company. This new software is capable of helping carriers achieve all three of these objectives. It will deliver the level of service *and* network control that carriers need as they move forward with their strategies to deliver premium, profitable, data connectivity services. By gaining control of these strategic assets, they will be better able to truly regain control of their businesses.

Availability Measurement

Craig Tysdal
President and Chief Executive Officer
NetSolve, Inc.

The networking community has understood for years that a network must exceed a certain availability, and reliability, standard to have business value. That explains why the concept developed that very cheap bandwidth has little value to a business because it either breaks down frequently or takes too long to repair. This de facto standard was applied with more anecdotal evidence than real science. Still, it became an easy excuse for an understaffed networking organization to shrug and exclaim, "We have a choice here—we can fix the problem, or we can document the outage." Such excuses are inexcusable today. With networks being the backbone of today's business, measurement systems are critical to effective network management.

The Internet's growth, the dependence of businesses on a network for revenue and customer satisfaction, and the advent of new technologies will increase the requirement that enterprises better understand and measure availability. In particular, technologies such as voice over Internet protocol (VoIP), digital subscriber line (DSL), and cable modems represent opportunities for business to become more cost-effective, provided that the availability of these subsystems exceeds business requirements. This paper discusses the general concepts of availability measurement, and it offers five critical questions that availability-measurement systems should address.

1. What availability metrics are meaningful?
2. What is the availability of the network to our internal and external customers?
3. Are vendors meeting the availability expectations agreed upon?
4. What factors, including people, tools, training, process, vendors, and subsystems, must be improved to increase availability?
5. How does my network availability compare with that of my peers/competitors?

These questions are meaningful when examined from multiple perspectives, including those of the enterprise, networking vendors, and management service providers. This paper's intent is not to examine the technical issues, but rather to understand the completeness of a solution from a business perspective.

What Availability Metrics Are Meaningful?

For availability metrics to become a business tool, it is crucial to understand and define the set of metrics that is truly important. This will avoid at least two common mistakes. First, that a lack of focus and specificity will yield too much data. Too much data will elongate the analysis cycle and often prevent or rush the most important step: corrective action. Second, that an incorrect choice or focus on a metric may produce counter-productive behavior. Focusing solely on mean time to repair (MTTR), for example—outside the context of availability and number of incidents per device/application—could mean that success or a better metric is achieved by a large number of small outage situations. Therefore, chronic sites may exist that are not obvious by examining the MTTR metric without additional data. Consequently, a frame-relay provider may achieve a significant competitive advantage in marketing a lower MTTR number when, actually, the network is less stable, and the MTTR metric is "enabled" by excessive PVC/Port bounces or cleared during test incidents.

NetSolve focuses on six key metrics: managed availability, unmanaged availability, proactive notification, number of incidents per subsystem per month, mean time to identify problems, and MTTR. Rather than define each of these metrics in detail, we will outline the business thinking and intent integrated into their definitions.

Managed Availability

The concept of a "managed" device, subsystem, or application is familiar to a management service provider. It is the set of components necessary for operating and maintaining. This concept also applies to the enterprise, with good definitions setting expectations. Great customer service is impossible without solid expectations.

Once a clear understanding of responsibility has been set, the concept of availability depends on simple arithmetic and the definition of downtime. The guiding principle to defining downtime is simply whether the application or the network is available to the customer. Typically, from an industry perspective, downtime does not include time for scheduled maintenance or time associated with the failure of a primary link as long as a redundant link is operating correctly. Such times should be reported, trended, and managed. For example, redundant links in a wide-area network (WAN) may have significant performance and cost implications important to the business and user experience. Scheduled maintenance can increase to be more proactive in your management techniques. This is especially true when diagnosing a network condition that has not completely failed. Many troubleshooting techniques and processes center around a "hard down" condition. Additionally, scheduled maintenance frequently occurs when there will be minimal business impact (if such a thing exists in your business). So, although these times typically are excluded from a "managed" metric, including them in an availability calculation is meaningful to ascertain trends and to manage.

Finally, the dynamics of defining downtime changes, depending on whether you are proactive or reactive in reporting it. In a reactive scenario, when capturing downtime using a trouble-ticket system, the downtime typically starts when a user notifies a help desk or vendor. This time may significantly differ from the time when the actual failure occurred. Ignoring the somewhat-academic argument about whether the network is down if no user or machine requires its use, the reality is that proactive management tends to capture more downtime and more incidents of downtime. Therefore, comparisons of availability calculations require additional information.

Unmanaged Availability

Unmanaged availability metrics attempt to define the entire user experience regardless of whether the management service provider, vendor, or internal support organization operates and supports it. We maintain that any measurement system must create metrics that closely parallel the total customer experience. For example, if the network runs perfectly, but poorly trained mobile users struggle with correctly connecting or using it, the business purpose is thwarted and customer satisfaction significantly hurt. Defining the network-system components that accurately reflect the customer impression requires thought. Additionally, the natural inclination would be to increase complexity. But that should be avoided. Customer responses on their own experience tend to be fairly simple; "it works great," "it stinks," "it works better," "it operates the same," or "it's worse than last month."

Proactive Notification

While this metric has no real implication to network availability, it does save time for customers and users—and improves their mental health. The faster a customer or user understands that there's a problem with the network and when it is expected to be fixed, the quicker that business contingency plans can be made. In addition, in those instances when responsibility for information technology (IT) is divided, i.e., network/application or voice/data, proactive notification by the responsible party to the business partner can save troubleshooting time and, hence, money.

Number of Incidents per Subsystem per Month

This metric is useful to detect situations where there are numerous interruptions of service, but they require little time to repair, known as MTTR. Thus, user dissatisfaction may not be high, since repairs are made quickly. Converting this metric into mean time between failure (MTBF) will allow some consistency of language in discussions with software and hardware vendors. Capturing this metric has an additional benefit: Networking equipment and services providers will devote different troubleshooting processes and resources to situations that have a documented history of repeated failure. Finally, this metric can help us to understand the quality of our network-design and configuration-management practices.

Mean Time to Identify Problems

The time required to identify the cause of a failure is only a component of MTTR and, in and of itself, is not that meaningful in a business sense. Nonetheless, this is an important element to managed service providers and internal IT organizations wanting to measure effectiveness, efficiency, and return on investment (ROI) of people, tools, and processes.

Mean Time to Repair

MTTR is the mean time interval between a failure's detection and the time it takes to repair it and restore service to the user. Calculations of this metric and its importance depend on understanding what is included in downtime and incident-failure analysis. Histograms of this metric as well as breakouts of this information to appropriate subsystem and supplier are very important.

The highlighted metrics are not the inclusive list that we internally monitor and/or publish for customers. They are an excellent place to begin the discussion about availability, and they prepare us to move to the second important question.

What Is the Availability of Our Network to Internal and External Customers?

Our customers tell us that different applications and different users require different levels of availability. Additionally, designing in high levels of availability in networks traditionally increases cost. Finally, when considering customer service, it helps to look at availability the same way that a customer might. We should allow customers to determine many of the business rules that define an availability calculation. For example, a customer might want to look at only one site or device in the last week or two. This same customer may believe that something less than an around-the-clock, seven-day-a week calendar is critical to truly understand the impact of associated outages on the business. Under this approach, the customer is free to make commitments to internal and external customers (read service-level agreements [SLAs], if you will) and maintain a flexible reporting system for appropriate measurement of those commitments. But can vendors then tailor their respective SLAs so that SLAs made from the customer to its customers conveniently flow down? This is the preferred situation.

These same market dynamics have prompted some software vendors to provide tools that correlate applications with users, thus defining availability from an application perspective. A high degree of automation is required in this approach to ensure that the underlying integrity of the database yields meaningful results that are cost-effective to gather and maintain.

Are My Vendors Meeting the Agreed-Upon Availability Expectations?

The availability system, to the highest degree possible, must reflect the availability and responsiveness of the responsible vendors. The Internet's value proposition for virtual private networks (VPNs) degrades significantly if the availability in the local loop or in a multicarrier long-haul environment is significantly more fragile than such private network technologies as private line or frame relay. Similarly, robust implementation of VoIP should not be tainted in any way if on-or-off net-wide area network links fail.

To answer this question, the enterprise or management service provider will measure certain important metrics such as downtime at higher degrees of granularity. For example, a component of downtime for an MTTR calculation when a network hardware failure occurs is the time it takes the vendor to respond to a dispatch. Clearly, if a customer has purchased four-hour response time, then it is critical that the vendor arrives within four hours of notification. It is assumed, of course, that the vendor has the correct replacement part and that it works. Availability systems need to measure this responsiveness as well as the underlying assumption of effective resolution.

What Factors—Including Training, Process, Vendors, and Subsystems—Do I Need to Improve to Increase Availability?

The requirement for measuring availability to help answer this question is the ability to trap, trend, and report on elements of availability with a high degree of specificity. Key to this are the error codes and the accurate reporting of the same. This process has some natural enemies. When critical systems are down and multiple vendors are working on a problem, it is not unusual that more than one component will be analyzed or changed simultaneously. If this works, then it is difficult to determine the problem's true cause. Subsequent failure analysis does not always unwind this problem and, given the repair-and-replacement cost of networking components today, many times it is not even seriously considered. Accurate reporting also assumes integrity among customer employees, vendors, and management service providers. In our business, we believe the customer should have control over their network even though they have hired us to operate it. As a result, we many times will share password control of configuration management. We have found that the use of authorization systems, nightly archives, and exception reporting is invaluable to determining where our processes and individual training require improvement.

How Does My Network Availability Compare with Peers/Competitors?

Once a customer has reported and trended availability, the obvious question becomes: How good is it? Is meeting our internal and vendor SLAs sufficient? Is a competitor's IT organization beating our brains out? The advantage of a management service provider with a solid availability reporting system is access to statistically significant amounts of comparative data. This data—when available in easy to understand formats—will answer the central question. In addition, it potentially becomes very useful for vendor management and evaluation. Are there availability costs to a free Internet service? What are the comparative results of different DSL vendors or different broadband access technologies? The networking industry has embraced published results on lab tests and some Internet-related operations data. Significant opportunity exists to increase this information and augment or replace marketing claims with operational reality.

Conclusion

The availability of our network last month was 99.98%. This is a simple statement. The purpose of this white paper is to provide a perspective on defining a robust availability reporting system. If properly designed, understood, and delivered, this reporting system will help enterprises to answer simple questions so they, too, can define their network availability just as easily.

What Does Redundancy Mean, Really...

James Wilson
Product Marketing and Program Manager
General Bandwidth

The technology that powers the public switched telephone network (PSTN) has been designed, refined, and optimized over the last several decades. At the heart of PSTN reliability is redundancy. Redundancy in carrier-class networks, simply defined, is the ability of telecommunications equipment to continue operation without service outages in the case of failure of any component of the telecommunications equipment.

The convergence of voice and data networks is making it possible to deliver voice, data, and multimedia over a single broadband network, driving the transformation of the PSTN from a circuit-switched network to an end-to-end packet network dubbed the new voice infrastructure (NVI).

The challenge is to provide the same level of redundancy in the NVI that has come to be expected of the PSTN.

Redundancy Defined

Carrier-class redundancy is built into every facet of packet-based voice networking equipment that guarantees the same level of reliability in NVI equipment as that offered by existing PSTN equipment. For example, in voice over digital subscriber line (VoDSL) and voice over Internet protocol (VoIP) applications, products with carrier-class redundancy are able to maintain active calls and minimize the impact to calls in the process of being set up when a component of the equipment fails.

While important to service providers for every application, carrier-class redundancy is critically important for equipment delivering derived voice over broadband networks. Voice gateways, which provide the center point for linking the packet- and circuit-switched networks are particularly susceptible to any failure, making redundancy that much more critical. The loss and retransmission of voice packets directly affects voice quality. If voice packets arrive out of sequence or with some packets missing, the end result is garbled communication. Because VoDSL and VoIP gateways have to work as faultlessly as the PSTN, they must be held to a higher redundancy standard than data networks. As a further requirement, services enabled by voice over broadband gateways must be completely transparent to the end user in terms of availability, reliability, and quality. *Figure 1* summarizes data from a Bellcore study on factors affecting telephone network downtime.

Software, Hardware, and Environmental Conditions

Outages due to software, hardware, and environmental conditions contribute to 45 percent of telephone network downtime. In a carrier-class redundant system, the software design leverages the hardware redundancy model to minimize the impact to system availability during outage periods and maintenance activities. A distributed, modular architecture is typically employed to ensure reliable system operation. This architecture includes well-defined software components and interfaces that clearly delineate control and status functions. A well-defined software architecture greatly assists in the management of system resources and will quickly direct redundancy resources when necessary.

Hardware redundancy typically involves using one or more spare devices to compensate for a failed active component during normal system operation. Upon failure of the primary device, the secondary device assumes operation with no interruption in service. This combination of a primary device and a secondary device comprise the minimum set for a protection group. Redundant devices in voice gateway products include module, port, timing system, and power-system redundancy.

Module Redundancy

Modules, or cards, provide system control and input/output (I/O) functionality in voice gateway equipment. Each module must be redundant in order to guarantee carrier-class redundancy. There are different redundancy techniques used depending on the type of module.

1:1

1:1 redundancy is typically used for system controller modules and I/O interfaces. In a 1:1 redundancy scheme, one primary module and one secondary module comprise one protection group. There are three service state models for 1:1 protection groups. VoDSL and VoIP gateways must employ Active/Hot Standby or Active/Active Standby to be considered carrier-class products.

FIGURE 1

Factors Affecting Telephone Network Downtime

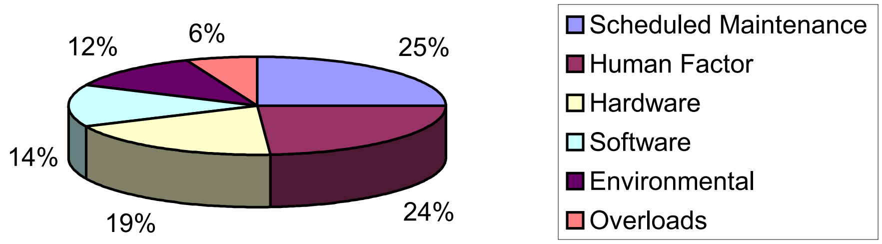

With Active/Hot Standby, both static and dynamic data is maintained in the Hot Standby secondary module. The Hot Standby card queues all activity and is in synchronicity with the primary card but does not act on requests or transactions. If the primary module fails, the secondary module can quickly assume traffic responsibilities because of its dynamic data synchronicity with the primary module. Active/Hot Standby facilitates quicker switchover times than Active/Cold Standby and is used in carrier-class redundant systems.

In Active/Active Standby, both the primary and secondary modules are active and both act on requests and transactions. The active card is designated the primary card and transmits data, while the standby card processes but does not transmit data. This approach minimizes switchover times and is used in carrier-class redundant systems.

1:N

1:N redundancy is used to maximize the number of in-service I/O interfaces because one secondary module can back up N primary modules in a protection group. This is especially relevant in trunking applications, as duplication of interface modules can be expensive and 1:N redundancy minimizes the cost of leasing spare lines.

In 1:N configurations, the spare module cannot be fully configured because its exact configuration is not completely known until one of the primary modules in the protection group fails. At that time, the secondary card is updated and begins to transmit traffic. For this reason, 1:N systems typically are Active/Cold Standby.

Hot Swappability

Another carrier-class requirement involving module redundancy is hot swappability, which is the ability to replace defective modules with non-defective modules while not impacting active traffic through the system. Modules that are hot swappable must also have the ability to automatically download their proper configuration when plugged into the slot.

Port Redundancy

In some instances, it is important to provide redundancy at the port level in addition to the card level. This is especially critical in situations where a large percentage of traffic is flowing over a handful of ports and loss of a single port can have a substantially negative impact on overall system availability. *Figure 2* details an example of this type of redundancy.

In this diagram a digital signal (DS)–3 Y-cable provides port redundancy. Both transmit signals are connected through a Y-cable to a transmit line, while both receive signals are connected through a Y-cable to a receive line. 75-ohm splitters determine which port is primary and which is secondary. Upon failure of the primary port, the secondary port begins transmission with no signal loss. Port redundancy is critical for VoDSL and VoIP gateways to provide a degree of fault tolerance above and beyond module-level redundancy.

Timing System Redundancy

To ensure fault tolerance and flexibility, multiple timing sources must be available for carrier-class voice gateways. This must include both internally generated timing sources as well as externally generated timing sources. Common timing sources include BITS and CC clocks. Upon failure of all internal and external timing sources, a Stratum Holdover clock must be available on carrier-class voice gateways.

Power System Redundancy

Some networking equipment includes internal power supplies. This equipment cannot be considered carrier class, as it introduces another potential point of failure that could adversely impact the system's overall mean time between failure (MTBF). Carrier-class VoDSL and VoIP gateways typically use a remote redundant power source and two 48-volt direct current (DC) redundant power inputs. In the case of failure of one power input, the second input is available to ensure the equipment remains powered.

Scheduled Maintenance

Scheduled maintenance (25 percent) represents a significant contributor to network downtime. Equipment design not only influences how often equipment must be serviced, in terms of reliability, but how the scheduled maintenance procedure actually occurs. For voice over broadband gateways, system availability during scheduled maintenance is especially critical. Carrier-class products employ split mode operation, a feature that allows equipment to be upgraded in the field without service interruption. Split mode is typically used to facilitate new software loads or hardware replacement. While a software or hardware update is made on one side of the split, normal transactions occur on the other side of the split. Split mode allows carrier-class gateways to continue processing active calls and set up new call connections during the maintenance process. Because of this capability, split mode is a necessary component for carrier-class voice gateways.

FIGURE 2

DS–3 Y-Cable Redundant Device Configuration

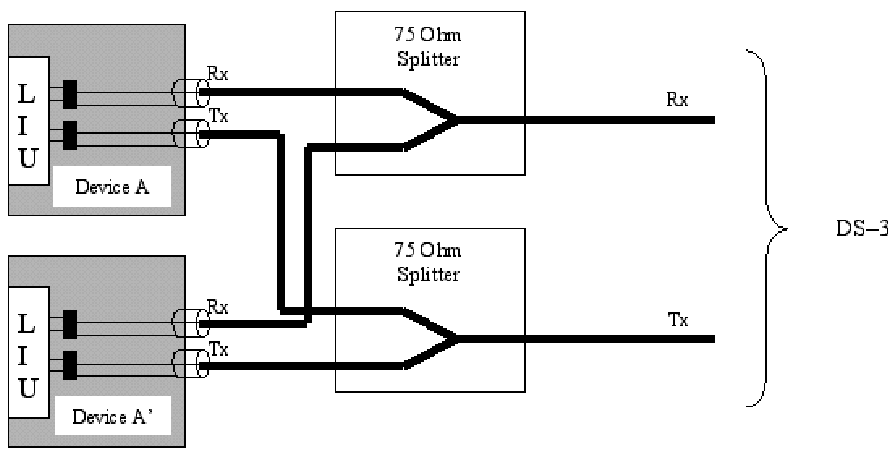

Human Factors

Human factors (24 percent) are another major cause of downtime. In addition to split mode, the ability to maintain equipment in a straightforward manner influences carrier-class service. Every field replaceable unit on a gateway should be clearly labeled to minimize confusion during installation and upgrades. Gateway design should be modular enough that removal of unrelated hardware or cables is minimized. Carrier-class gateway products typically employ a midplane architecture that allows service technicians to replace faulty components while not affecting the remainder of the system.

Another important human factor consideration is the design and appearance of the gateway element-management system (EMS). The EMS should be clearly organized to allow for easy fault, performance, and configuration management of the gateway. An intuitive point-and-click graphical user interface (GUI) is a necessary feature. Another important feature of an EMS is its event notification process. All equipment state changes should be propagated from the gateway to the EMS and clearly identify to the user what is happening on the gateway under management in real time. Diagnostic tools should be available to monitor and measure gateway performance from the EMS. Further, the EMS should allow the user the same degree of management control remotely as a user on site with the equipment.

Traffic Overloads

Traffic overloads contribute another 6 percent to network downtime. Carrier-class gateways have the ability to directly control packetized voice traffic. In VoDSL applications, carrier-class gateways are able to shape and police traffic on multiple levels: virtual path (VP), virtual circuit (VC), and channel identifier (CID). This gives users the maximum amount of flexibility to manage traffic across their voice over broadband networks.

In voice gateways, system bus capacity can also impact traffic conditions. Carrier-class gateways are based on asynchronous transfer mode (ATM) bus architectures that provide in excess of 10 gigabits per second (Gbps) backplane/midplane capacity. These ATM–based designs provide sufficient excess bandwidth that the system bus does not become a bottleneck, even in high-traffic conditions. This design is superior to compact peripheral component interconnect (PCI)–based architectures, which are limited to 1 Gbps or less capacity and become bottlenecks in high-traffic situations. Another limitation of compact PCI–based voice gateways is that they require special PCI bridge circuitry to extend beyond eight slots, which introduces another potential point of failure in the system.

Conclusion

Through more than 100 years of development, the PSTN in the United States has been the most reliable network in the world. Unlike the Internet and corporate local-area networks (LANs), where network outages invariably occur from time to time, phone service has to be error-free. Reliable phone service is of paramount importance because telephony service plays such a critical role in the everyday life of business and residential subscribers.

Carrier-class redundancy in VoDSL and VoIP gateways plays a critical role in ensuring the system reliability and availability that enables packet voice services to meet the requirements of traditional voice services.

QoS Metrics for SLA Management

Eric Yam
Chief Technology Officer
ECtel

The telecommunications industry is a dynamically changing environment. New technologies, such as Internet protocol (IP), wireless, and digital subscriber lines (DSL), are rapidly emerging, and new servers and new business models are being introduced. This paper discusses quality of service (QoS) for service-level agreement (SLA) management for new services from the perspective of service management and business management.

New QoS Challenges

QoS was traditionally an operations and engineering tool used by QoS experts for a relatively static, well-established connection-based network environment. Today it has become a business management tool with immediate commercial implications.

Meeting Emerging Telecommunications Business Transactions

QoS is now being used to formulate and enforce the compliance of SLAs. Telecommunications capacity is being traded among alternative carriers as if it were a stock market commodity, and price is determined by volume and quality. Best-value routing (BVR) is replacing the traditional least-cost routing by using a value criteria, which is the cost weighted by quality.

Requirements for SLAs

The market should be allowed to provide different service grades at different price levels, and QoS metrics should be used for the formulation of SLAs and subsequent compliance confirmation.

Industrially standardized QoS methodology is needed to characterize true customer-perceived QoS, and that should be applicable to different services and technologies. In addition, the tool should be an automated, hands-free quality measurement of live traffic, supported by near-real-time reporting and information distribution.

Also needed are short- and long-term quality statistics, which can be measured against target QoS benchmarks, and this requires metrics.

QoS Metrics

Metrics is an index to represent true customer perception of call quality, and it is based on well-researched and well-tested models. It is used for quantifying abstract quality attributes in simple numbers that are easy to interpret and use by experts and amateurs alike. Metrics is a standard for industry-wide use that enables everyone to speak the same language.

Call-Quality Metrics

Telephone call quality is based on four factors: the call-connection success rate, call-connection delay, conversation and voice quality, and fax transmission quality.

Call-Connection Success Rate

The call-connection success rate includes the answer-seizure ratio (ASR) and the call-completion ratio (CCR). It is based on the number of attempts made before a connection is completed.

Call-Connection Delay

The call-connection delay includes the post-dial delay (PDD) and the call setup time—i.e., how long it takes to get a successful connection. This is an important measurement to the service provider because call-connection delay is not billable time. It is important to the user because waiting can lead to frustration.

Conversation and Voice Quality

The conversation and voice quality is a value measured by the call clarity index (CCI), the perceptual analysis measurement system (PAMS), and the perceptual evaluation of speech quality (PESQ). The CCI is now the new International Telecommunications Union–Telecommunication Standardization Sector (ITU–T) Rec. P.562, and the new standard for speech quality is the ITU–T P.862.

Fax Transmission Quality

The fax transmission quality includes the number of page sends, the data rate, and image quality. The new standard for fax transmission quality is ITU–T Rec. E.458.

Call-Quality Measurement

There are two basic approaches to measuring call quality. One method is to measure live traffic, and the other is to generate test calls. A hybrid of these two approaches is also possible.

Live-Traffic Measurement

Measuring live traffic is nonintrusive. It is a high-volume per-call measurement per the ITU–T Rec. P.561, which could amount to measuring tens of millions of call records a day and storing up to a year's worth of data.

With the information in a call-record database, statistics can be generated per carrier or per destination over a selected time interval, such as hourly, daily, or weekly, with centralized control and scheduling. It will also provide a macroscopic view of network health.

The high-level metrics involved in live-traffic measurement include the call connectivity and PDD, CCI, or ITU–T Rec. P.562, and the fax performance index.

Test Calls
Generating test calls is intrusive. It requires centralized control of strategically located remote test units (RTU) to make test calls via the network being tested. The microscopic view of the system that intrusive testing provides is complementary to the macroscopic view provided by the nonintrusive method.

The test call measures call connectivity and setup delay as well as user-perceived quality for voice over IP (VoIP), using the psycho-acoustic model PAMS/PESQ. The test call also has to measure IP network parameters.

The Hybrid Approach
The hybrid approach to call-quality measurement makes use of both nonintrusive and intrusive methods. An RTU is strategically located for the generation of intrusive test calls, and the traffic center houses the nonintrusive live-traffic measurement devices (see *Figure 1*).

Mapping of Network Parameters into User-Perceived Quality
Call-quality measurement results in a high volume of data. Measurements must be taken for speech level or loudness, noise, echo loss and delay, voice compression effects, and absolute delay, which includes transmission and processing. The latency, delay variation or jitter, packet loss percentage, errored packet percentage, and the signal enhancement scheme must also be measured.

To use all these measurements for service-level management, they need to be mapped into user-perceived quality. The user-perceived mean opinion score (MOS) ranges from five to one, where five is excellent and one is bad. For each call there must actually be two MOSs, one to represent each direction of the call.

Average CCI for All Calls
The average CCI for many carriers falls between four, which is good, and three, which is fair. *Figure 2* shows the actual measurements of live traffic, using the CCI for 14 different carriers, labeled A through N. There is some variation among these carriers, but not much.

Average CCI for Bad Calls
Looking at only the bottom 20 percent of bad calls for the same carriers reveals a large amount of variation (see *Figure 3*). The CCI of carrier B has dropped from about 3.3 (better than fair) to below 2.5 (midway between fair and poor). That of carrier F has dropped from almost 3.5 (midway between good and fair) to about 2.9 (less than fair but not as poor as carrier B).

Customers who placed the poor-quality calls shown for carriers B and F would most likely complain and ask for a refund.

The CCR
The CCR can be measured hourly for a 24-hour period and then charted to show the distribution of completed, incomplete, and rejected calls. *Figures 4* and *5* illustrate the hourly

FIGURE 1

Hybrid Approach Combining Nonintrusive and Intrusive Methods

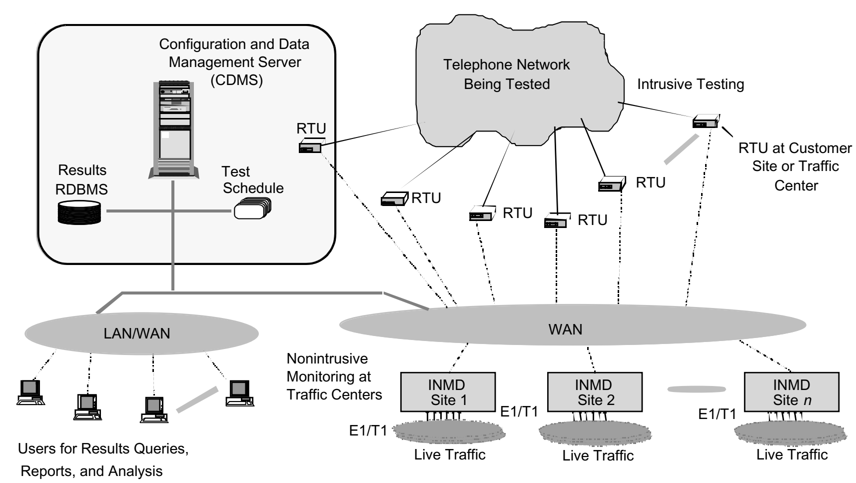

Key
INMD in-service, nonintrusive measurement device
LAN local-area network
RDBMS relational database management system
WAN wide-area network

FIGURE 2

Average CCI Values by Carrier

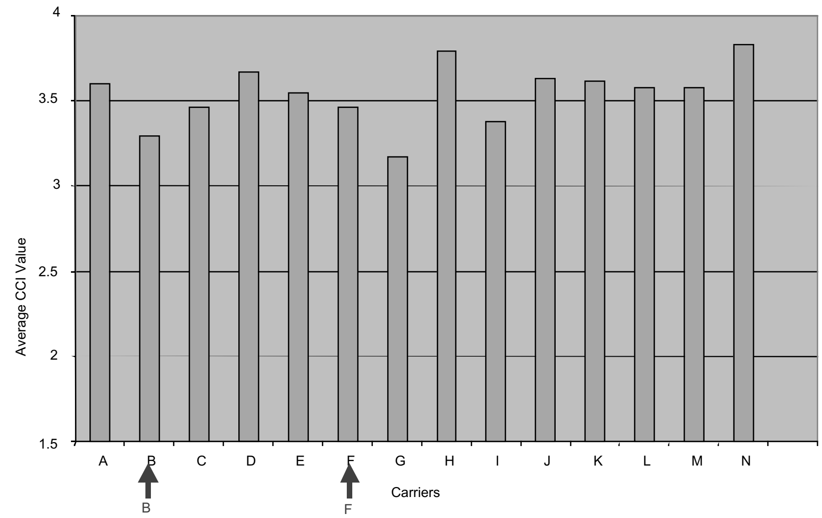

breakdown of calls for a single carrier on a single day, first for phone calls and then for fax calls. Based on the ratio of completed calls to incomplete and rejected calls, the shown carrier's record is not very good.

Voice-Quality Benchmarking

Intrusive tests, especially for listening quality (LQ), would be even more revealing. A comparison of LQ MOS, measured by the PAMS intrusive test, was done for a local public switched telephone network (PSTN), an international PSTN, and a VoIP phone-card service (see *Figure 6*).

The international carrier in *Figure 6* did slightly worse than the local carrier, but the quality of both was nonetheless very consistent and good. Not only is the average VoIP quality poorer than the international carrier, there are also some very bad calls.

Industrial Standardization

The ITU–T study group is doing a lot of work in the area of industrial standardization, and the QoS of VoIP is being worked on. The QoS Development Group (QSDG) and the ITU–T are very active in this area, as is the Internet Engineering Task Force (IETF).

ETSI TIPHON QoS Classes

The European Telecommunications Standards Institute (ETSI) Telecommunications and Internet Protocol Harmonization over Networks (TIPHON) considers the call

FIGURE 3

Worst 20 Percent Average CCI Values by Carrier

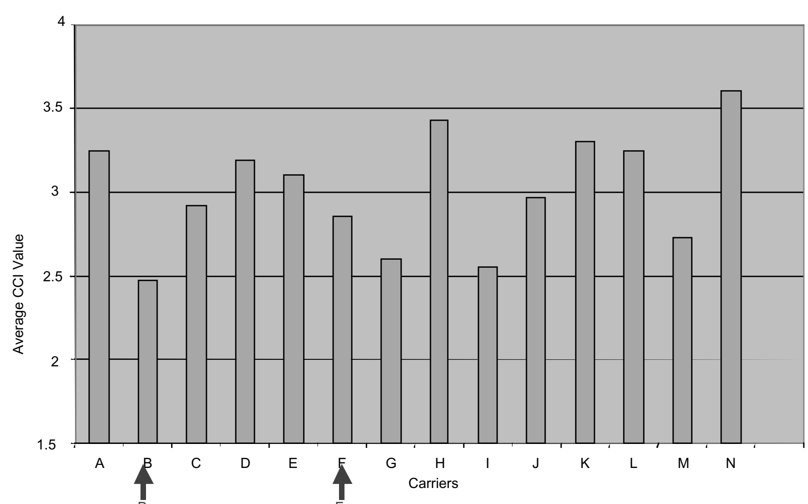

FIGURE 4

CCR for 24 Hours

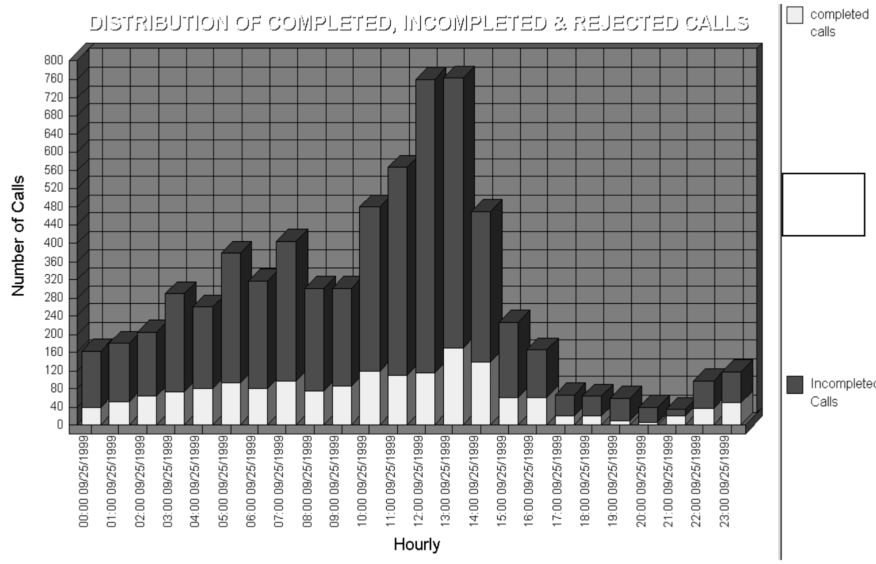

as a service and defines four QoS classes, per DTS-101-512, March 2000:

- 4 = Best
- 3 = High
- 2 = Medium
- 1 = Best effort

These classes are expressed in terms of several measurable parameters, including overall transmission quality, speech quality, end-to-end delay, and call setup time.

IPDR.org

IPDR.org is an industrial organization for facilitating practical implementation of interchange of usage data among service elements of IP–based services, which include the following:

- Session-based calls (versus call-based records)
- Content-based billing (versus connection-based billing)
  - QoS–based billing (versus best effort)
  - Multiple services such as VoIP, video on demand (VOD), e-mail, wireless application protocol (WAP), and multimedia conferencing

FIGURE 5

Fax Success Rate for 19 Hours

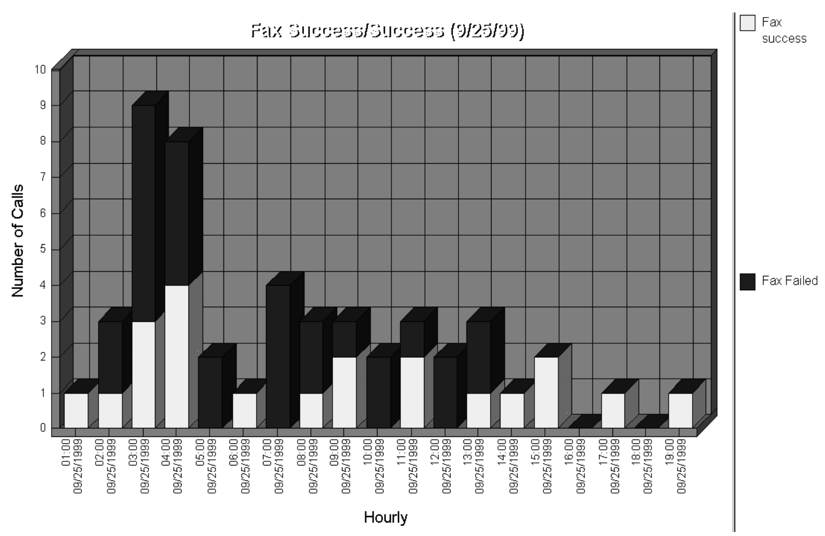

FIGURE 6

Voice-Quality Benchmarking

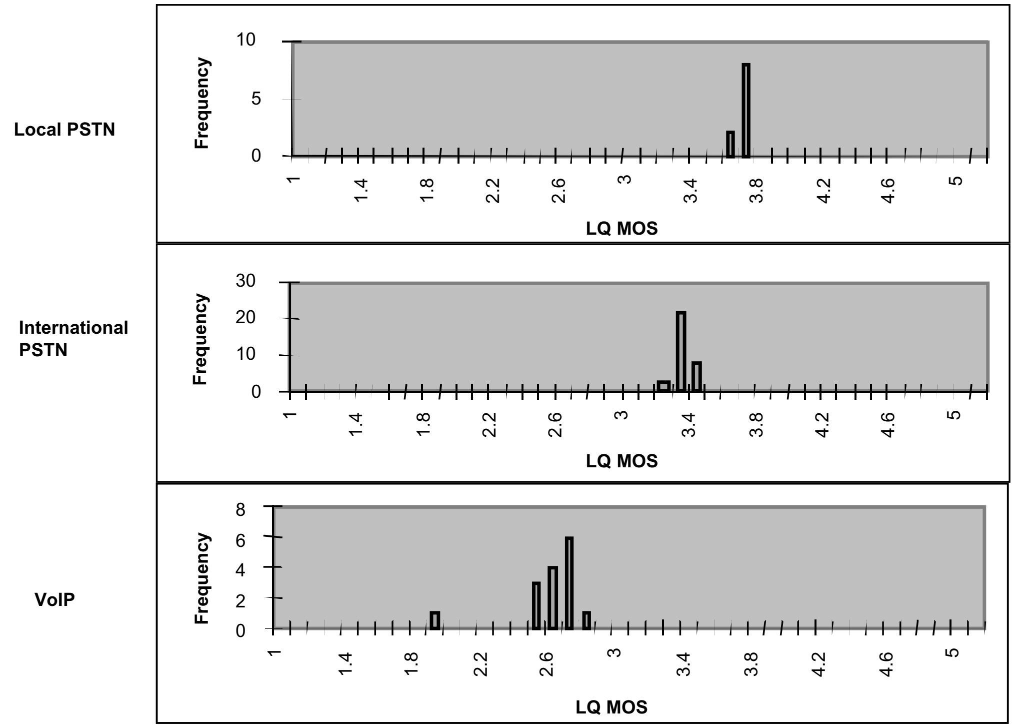

Network data management–usage (NDM–U) version 2.5, scheduled for release in first quarter 2001, will include VoIP and VOD IP data record (IPDR) specifications with QoS attributes. IPDR work will be essential for SLA management for IP–based services.

Looking Ahead

Fundamental QoS tool sets will be available for SLA management of emerging services, such as IP–based voice service. New tools will be needed for specific new services, such as a multimedia opinion model.

It will be an evolutionary and collaborative process to apply the tools to fit different business scenarios and industry standardization efforts by the ITU–T, QSDG, TIPHON, and IPDR.org.

Acronym Guide

2B1Q	two binary, one quaternary
2G	second generation
3DES	triple data encryption standard
3G	third generation [also G3]
3GPP	third-generation partnership project
3R	regeneration, reshaping, and retraining
4B3T	four binary, three ternary
4F/BDPR	four-fiber bidirectional dedicated protection ring
4F/BSPR	four-fiber bidirectional shared protection ring
AAA	authentication, authorization, and accounting
AAL–[x]	ATM adaptation [layer x]
ABC	activity-based costing
ABR	available bit rate
AC	alternating current OR authentication code
ACD	automatic call distributor
ACH	automated clearinghouse
ACL	access control list
ACLEP	adaptive code excited linear prediction
ACM	address complete message
ACR	alternate carrier routing OR anonymous call rejection
ADM	add/drop multiplexer OR asymmetric digital multiplexer
ADPCM	adaptive differential pulse code modulation
ADS	add/drop switch
ADSI	analog display services interface
ADSL	asymmetric digital subscriber line
AFE	analog front end
AGW	agent gateway
AIM	advanced intelligent messaging
AIN	advanced intelligent network
ALI	automatic location identification
AM	amplitude modulation
AMA	automatic messaging account
AMI	alternate mark inversion
AMPS	advanced mobile phone service
AN	access network
ANI	automatic number identification
ANM	answer message
ANSI	American National Standards Institute
ANX	Automotive Network Exchange
AOL	America Online
AON	all-optical network
AP	access point OR access provider
APC	automatic power control
API	application programming interface
APON	ATM passive optical network
APS	automatic protection switching
ARCNET	attached resource computer network
ARI	assist request instruction
ARM	asynchronous response mode
ARMS	authentication, rating, mediation, and settlement
ARP	address resolution protocol
ARPANET	Advanced Research Projects Agency Network
ARPU	average revenue per customer
A-Rx	analog receiver
AS	application server OR autonomous system
ASAM	ATM subscriber access multiplexer
ASC	Accredited Standards Committee
ASCII	American Standard Code for Information Interchange
ASE	amplified spontaneous emission
ASIC	application-specific integrated circuit
ASIP	application-specific instruction processor
ASON	automatically switched optical network
ASP	application service provider
ASR	access service request OR answer-seizure rate OR automatic service request OR automatic speech recognition
ASSP	application-specific standard part
ASTN	automatically switched transport network OR analog switched telephone network
ATC	automatic temperature control
ATIS	Alliance for Telecommunications Industry Solutions
ATM	asynchronous transfer mode OR automated teller machine
ATMF	ATM Forum
ATP	analog twisted pair
ATU-C	ADSL transmission unit-CO
ATU-R	ADSL transmission unit-remote
A-Tx	analog transceiver
AUI	attachment unit interface
AVI	audio video interleaved
AWG	American Wire Gauge OR arrayed waveguide grating
AYUTOS	as-yet-un-thought-of services
B2B	business-to-business
B2C	business-to-consumer
BAF	Bellcore automatic message accounting format
B-Box	breakout box
BBS	bulletin board system
BCSM	basic call state model
BDCS	broadband digital cross-connect system
BDPR	bidirectional dedicated protection ring
BE	border element
BER	bit-error rate
BERT	bit error–rate test
BGP	border gateway protocol
BH	busy hour
BHCA	busy hour call attempt
BI	bit rate independent
BICC	bearer independent call control
BID	bit rate identification
BIP	bit interactive parity
B–ISDN	broadband ISDN
BLEC	broadband local-exchange carrier OR building local-exchange carrier
BLES	broadband loop emulation services
BLSR	bidirectional line-switched ring
BML	business management layer
BOC	Bell operating company

BOF	business operations framework
BOND	back-office network development
BOSS	broadband operating system software
BPON	broadband passive optical network
BPSK	binary phase shift keying
B–RAS	broadband–remote access server
BRI	basic rate interface
BSA	business services architecture
BSPR	bidirectional shared protection ring
BSS	base-station system OR business support system
BTS	base transceiver station
BVR	best-value routing
BW	bandwidth
CA	call agent
CAC	call admission control OR carrier access code OR connection admission control
CAD	computer-aided design
CAGR	compound annual growth rate
CALEA	Communications Assistance for Law Enforcement Act
CAM	computer-aided manufacture
CAMEL	customized application of mobile enhanced logic
CAP	competitive access provider OR carrierless amplitude and phase modulation
CAPEX	capital expenditures
CAR	committed access rate
CARE	customer account record exchange
CAS	channel-associated signaling OR communications applications specification
CAT	conditional access table OR computer aided telephony
CATV	cable television
C-band	conventional band
CBDS	connectionless broadband data service
CBR	constant bit rate
CBT	core-based tree
CC	control component
CCB	customer care and billing
CCF	call-control function
CCI	call clarity index
CCITT	Consultative Committee on International Telegraphy and Telephony
CCK	complementary code keying
CCR	call-completion ratio
CCS	common channel signaling
CD	chromatic dispersion OR compact disc
cDCF	conventional dispersion compensation fiber
CDD	content delivery and distribution
CDDI	copper-distributed data interface
CDMA	code division multiple access
CDMP	cellular digital messaging protocol
CDMS	configuration and data management server
CDN	control directory number
CDPD	cellular digital packet data
CDR	call detail record OR clock and data recovery
CD–ROM	compact disc–read-only memory
CDWDM	coarse wavelength division multiplexing
CE	customer edge
CEI	comparable efficient interface

CEO	chief executive officer
CER	customer edge router
CES	circuit emulation service
CERT	computer emergency response team
CES	circuit emulation service
CESID	caller emergency service identification
CEV	controlled environment vault
CFB/NA	call forward busy/not available
CFO	chief financial officer
CFP	contention free period
CGI	common gateway interface
CHCS	composite health care system
CHN	centralized hierarchical network
C–HTML	compressed HTML
CIC	circuit identification code
CID	caller identification
CIMD2	computer interface message distribution 2
CIM	common information model
CIO	chief information officer
CIP	classical IP over ATM
CIR	committed information rate OR Communications Industry Researchers
CIT	computer integrated telephone
CLASS	custom local-area signaling services
CLE	customer-located equipment
CLEC	competitive local-exchange carrier
CLI	command line interface
CLLI	common language location identifier
CLR	circuit layout record
CM	cable modem
CM&B	customer management and billing
CMIP	common management information protocol
CMISE	common management information service element
CMOS	complementary metal oxide semiconductor
CMRS	commercial mobile radio service
CMTS	cable modem termination system
CNAM	calling name (also defined as “caller identification with name” and simply “caller identification”)
CNAP	CNAM presentation
CO	central office
CODEC	coder-decoder
COI	community of interest
COO	chief operations officer
COPS	common open policy service
CORBA	common object request broker architecture
CORE	council of registrars
CoS	class of service
COT	central office terminal
COTS	commercial off-the-shelf
COW	cellsite on wheels
CP	connection point
CPAS	cellular priority access service
CPC	calling-party category (also calling-party control OR calling-party connected)
CPE	customer-premises equipment
CPI	continual process improvement
CPL	call processing language
CPLD	complex programmable logic device
CPN	calling-party number
CPU	central processing unit
CR	constraint-based routing

CR–LDP	constraint-based routed–label distribution protocol
CRC	cyclic redundancy check OR cyclic redundancy code
CRM	customer-relationship management
CRS	capacity, reach, and speed
CRTP	compressed real-time transport protocol
CRV	call reference value
CS	client signal
CS–[x]	capability set [x]
CSA	carrier serving area
CSCE	converged service-creation and execution
CSCW	computer supported collaborative work
CS–IWF	control signal interworking function
CSF	critical success factors
CSM	customer-service manager
CSMA/CA	carrier sense multiple access with collision avoidance
CSMA/CD	carrier sense multiple access with collision detection
CSN	circuit-switched network
CSP	communications service provider OR content service provider
CSR	customer-service representative
CSU	channel service unit
CSV	circuit-switched voice
CT	computer telephony
CT–2	cordless telephony generation 2
CTI	computer telephony integration
CTIA	Cellular Telecommunications & Internet Association
CTM	Center for Telecommunications Management
CTO	chief technology officer
CUG	closed user group
CWD	centralized wavelength distribution
CWDM	coarse wavelength division multiplexing
CWIX	cable and wireless Internet exchange
DAC	digital access carrier
DACS	digital access cross-connect system
DAM	DECT authentication module
DAMA	demand assigned multiple access
DARPA	Defense Advanced Research Projects Agency
DAVIC	Digital Audio Video Council
dB	decibel(s)
dBrn	decibels above reference noise
DB	database
DBMS	database management system
DBS	direct broadcast satellite
DC	direct current
DCC	data communications channel
DCF	discounted cash flow OR dispersion compensation fiber
DCLEC	data competitive local-exchange carrier
DCM	dispersion compensation module
DCN	data communications network
DCOM	distributed component object model
DCS	digital cross-connect system OR distributed call signaling
DCT	discrete cosine transform
DDN	defense data network
DDS	dataphone digital service
DECT	Digital European Cordless Telecommunication
demarc	demarcation point
DEMS	digital electronic messaging service
DES	data encryption standard
DFB	distributed feedback
DFC	dedicated fiber/coax
DGD	differentiated group delay
DGFF	dynamic gain flattening filter
DHCP	dynamic host configuration protocol
DiffServ	differentiated services
DIN	digital information network
DIS	distributed interactive simulation
DITF	Disaster Information Task Force
DLC	digital loop carrier
DLCI	data-link connection identifier
DLE	digital loop electronics
DLEC	data local-exchange carrier
DLR	design layout report
DM	dense mode
DMD	dispersion management device
DMS	digital multiplex system
DMT	discrete multitone
DN	distinguished name
DNS	domain name server OR domain naming system
DOC	department of communications
DOCSIS	data over cable service interface specifications
DOD	Department of Defense
DOJ	Department of Justice
DoS	denial of service
DOS	disk operating system
DOSA	distributed open signaling architecture
DOT	Department of Transportation
DOW	day of week
DP	detection point
DPC	destination point code
DPE	distributed processing environment
DPT	dial pulse terminate
DQoS	dynamic quality of service
D-Rx	digital receiver
DS–[x]	digital signal [level x]
DSAA	DECT standard authentication algorithm
DSC	DECT standard cipher
DSCP	DiffServ code point
DSF	dispersion-shifted fiber
DSL	digital subscriber line [also xDSL]
DSLAM	digital subscriber line access multiplexer
DSLAS	DSL–ATM switch
DSP	digital signal processor OR digital service provider
DSS	decision support system
DSSS	direct sequence spread spectrum
DSU	data service unit OR digital service unit
DTMF	dual-tone multifrequency
DTV	digital television
D-Tx	digital transceiver
DVB	digital video broadcast
DVC	dynamic virtual circuit
DVD	digital video disc
DVMRP	distance vector multicast routing protocol
DVOD	digital video on demand
DVR	digital video recording
DWDM	dense wavelength division multiplexing
DXC	digital cross-connect
E911	enhanced 911

EAI	enterprise application integration
EBITDA	earnings before interest, taxes, depreciation, and amortization
EC	electronic commerce
ECD	echo-cancelled full-duplex
ECRM	echo canceller resource module
ECTF	Enterprise Computer Telephony Forum
EDA	electronic design automation
EDF	electronic distribution frame OR erbium-doped fiber
EDFA	erbium-doped fiber amplifier
EDGE	enhanced data rates for GSM evolution
EDI	electronic data interchange
EDSX	electronic digital signal cross-connect
EFM	Ethernet in the first mile
EFT	electronic funds transfer
EJB	enterprise Java beans
ELAN	emulated local-area network
ELEC	enterprise local-exchange carrier
EM	element manager
EMI	electromagnetic interference
EML	element-management layer
EMS	element-management system OR enterprise messaging server
ENUM	telephone number mapping
E–O	electrical-to-optical
EO	end office
EoA	Ethernet over ATM
EOC	embedded operations channel
EoVDSL	Ethernet over VDSL
EPD	early packet discard
EPON	Ethernet PON
EPROM	erasable programmable read-only memory
ERP	enterprise resource planning
ESCON	enterprise system connection
ESS	electronic switching system
ETC	establish temporary connection
EtherLEC	Ethernet local exchange carrier
ETL	extraction, transformation, and load
eTOM	e-telecom operations map
ETSI	European Telecommunications Standards Institute
EU	European Union
EXC	electrical cross-connect
EURESCOM	European Institute for Research and Strategic Studies in Telecommunications
FAB	fulfillment, assurance, and billing
FAQ	frequently asked question
FBG	fiber Bragg grating
FCAPS	fault, configuration, accounting, performance, and security
FCC	Federal Communications Commission
FCI	furnish charging information
FCIF	flexible computer-information format
FDA	Food and Drug Administration
FDD	frequency division duplex
FDDI	fiber distributed data interface
FDF	fiber distribution frame
FDM	frequency division multiplexing
FDMA	frequency division multiple access
FDS–1	fractional DS–1
FE	extended framing
FEC	forward error correction
FEPS	facility and equipment planning system
FEXT	far-end crosstalk
FHSS	frequency hopping spread spectrum
FICON	fiber connection
FITL	fiber-in-the-loop
FM	fault management OR frequency modulation
FOC	firm order confirmation
FOT	fiber-optic terminal
FOTS	fiber-optic transmission system
FP	Fabry-Perot [laser]
FPB	flex parameter block
FPGA	field programmable gate array
FPLMTS	future public land mobile telephone system
FPP	fast-packet processor
FR	frame relay
FRAD	frame-relay access device
FSAN	full-service access network
FSC	framework services component
FSN	full-service network
FT	fixed-radio termination
FT1	fractional T1
FTC	Federal trade Commission
FTE	full-time equivalent
FTP	file transfer protocol
FTP3	file transfer protocol 3
FTTB	fiber-to-the-building
FTTC	fiber to the curb
FTTCab	fiber-to-the-cabinet
FTTEx	fiber-to-the-exchange
FTTH	fiber-to-the-home
FTTN	fiber-to-the-neighborhood
FTTS	fiber-to-the-subscriber
FTTx	fiber-to-the-x
FWM	four-wave mixing
FX	foreign exchange
G3	[see 3G]
GA	genetic algorithm
Gb	gigabit
GbE	gigabit Ethernet [also GE]
GBIC	gigabit interface converter
Gbps	gigabits per second
GCRA	generic cell rate algorithm
GDIN	global disaster information network
GDMO	guidelines for the definition of managed objects
GE	[see GbE]
GEO	geosynchronous Earth orbit
GETS	government emergency telecommunications service
GFF	gain flattening filter
GFR	guaranteed frame rate
Ghz	gigahertz
GIF	graphics interface format
GIS	geographic information services
GKMP	group key management protocol
GMII	gigabit media independent interface
GMLC	gateway mobile location center
GMPCS	global mobile personal communications services
GMPLS	generalized MPLS
GNP	gross national product
GOCC	ground operations control center
GPIB	general-purpose interface bus
GPRS	general packet radio service

GPS	global positioning system
GR	generic requirement
GRASP	greedy randomized adaptive search procedure
GSA	Global Mobile Suppliers Association
GSM	Global System for Mobile Communications
GSMP	generic switch management protocol
GSR	gigabit switch router
GTT	global title translation
GUI	graphical user interface
GVD	group velocity dispersion
GW	gateway
HBO	Home Box Office
HCC	host call control
HD	home domain
HDLC	high-level data-link control
HDML	handheld device markup language
HDSL	high–bit rate DSL
HDT	host digital terminal
HDTV	high-definition television
HDVMRP	hierarchical distance vector multicast routing protocol
HEC	head error control OR header error check
HEPA	high-efficiency particulate arresting
HFC	hybrid fiber/coax
HIDS	host intrusion detection system
HLR	home location register
HN	home network
HOM	high-order mode
HomePNA	Home Phoneline Networking Alliance [also HomePNA2]
HomeRF	Home Radio Frequency Working Group
HPC	high probability of completion
HPO	high-performance option
HQ	headquarters
HSCSD	high-speed circuit-switched data
HSD	high-speed data
HSIA	high-speed Internet access
HSP	hosting service provider
HTML	hypertext markup language
HTTP	hypertext transfer protocol
HVAC	heating, ventilating, and air-conditioning
HW	hardware
IAD	integrated access device
IAM	initial address message
IAS	integrated access service OR Internet access server
IAST	integrated access, switching, and transport
IAT	inter-arrival time
IBC	integrated broadband communications
IC	integrated circuit
ICD	Internet call diversion
ICDR	Internet call detail record
ICL	intercell linking
ICMP	Internet control message protocol
ICP	integrated communications provider OR intelligent communications platform
ICS	integrated communications system
ICW	Internet call waiting
IDC	Internet data center OR International Data Corporation
IDE	integrated development environment
IDES	Internet data exchange system
IDF	intermediate distribution frame
IDL	interface definition language
IDLC	integrated digital loop carrier
IDS	intrusion detection system
IDSL	integrated services digital network DSL
IEC	International Electrotechnical Commission OR International Engineering Consortium
IEEE	Institute of Electrical and Electronics Engineers
I-ERP	integrated enterprise resource planning
IETF	Internet Engineering Task Force
IFITL	integrated [services over] fiber-in-the-loop
IFMA	International Facility Managers Association
IFMP	Ipsilon flow management protocol
IGMP	Internet group management protocol
IGP	interior gateway protocol
IGRP	interior gateway routing protocol
IGSP	independent gateway service provider
IHL	Internet header length
IIS	Internet Information Server
IKE	Internet key exchange
ILA	in-line amplifier
ILEC	incumbent local-exchange carrier
ILMI	interim link management interface
IM	instant messaging
IMA	inverse multiplexing over ATM
IMRP	Internet multicast routing protocol
IMSI	International Mobile Subscriber Identification
IMT	intermachine trunk OR International Mobile Telecommunications
IMTC	International Multimedia Teleconferencing Consortium
IN	intelligent network
INAP	intelligent network application part
INAP AU	INAP adaptation unit
INE	intelligent network element
InfoCom	information communication
INM	integrated network management
INMD	in-service, nonintrusive measurement device
INT	[point-to-point] interrupt
InterNIC	Internet Network Information Center
IntServ	integrated services
IOF	interoffice facility
IOS	intelligent optical switch
IP	Internet protocol
IPBX	Internet protocol private branch exchange
IPcoms	Internet protocol communications
IPDC	Internet protocol device control
IPDR	Internet protocol data record
IPe	intelligent peripheral
IPG	intelligent premises gateway
IPO	initial public offering OR Internet protocol over optical
IPoA	Internet protocol over ATM
IPQoS	Internet protocol quality of service
IPSec	Internet protocol security
IPTel	IP telephony
IPv6	Internet protocol version 6
IPX	Internet package exchange
IR	infrared
IRU	indefeasible right to user
IS	information service OR interim standard
IS-IS	intermediate system to intermediate system

ISA	industry standard architecture
ISAPI	Internet server application programmer interface
ISC	integrated service carrier OR International Softswitch Consortium
ISDF	integrated service development framework
ISDN	integrated services digital network
ISDN–BA	ISDN basic access
ISDN–PRA	ISDN primary rate access
ISEP	intelligent signaling endpoint
ISM	industrial, scientific, and medical OR integrated service manager
ISO	International Organization for Standardization
ISOS	integrated software on silicon
ISP	Internet service provider
ISUP	ISDN user part
ISV	independent software vendor
IT	information technology OR Internet telephony
ITSP	Internet telephony service provider
ITTP	information technology infrastructure library
ITU	International Telecommunication Union
ITU–T	ITU–Telecommunication Standardization Sector
ITV	Internet television
IVR	interactive voice response
IVRU	interactive voice-response unit
IWF	interworking function
IWU	interworking unit
IXC	interexchange carrier
JAIN	Java APIs for integrated networks
JCAT	Java coordination and transactions
JCC	JAIN call control
JDBC	Java database connectivity
JDMK	Java dynamic management kit
JMAPI	Java management application programming interface
JMX	Java management extension
JPEG	Joint Photographic Experts Group
JSCE	JAIN service-creation environment
JSIP	Java session initiation protocol
JSLEE	JAIN service logic execution environment
JTAPI	Java telephony application programming interface
JVM	Java virtual machine
kbps	kilobits per second
kHz	kilohertz
km	kilometer
L2F	Layer-2 forwarding
L2TP	Layer-2 tunneling protocol
LAC	L2TP access concentrator
LAI	location area identity
LAN	local-area network
LANE	local-area network emulation
LATA	local access and transport area
L-band	long band
LBS	location-based services
LC	local convergence
LCD	liquid crystal display
LCP	link control protocol
LCUG	Local Competition User Group
LD	laser diode OR long distance
LDAP	lightweight directory access protocol
LD–CELP	low delay–code excited linear prediction
LDP	label distribution protocol
LDS	local digital service
LE	line equipment OR local exchange
LEAF®	large-effective-area fiber
LEC	local-exchange carrier
LED	light-emitting diode
LEO	low Earth orbit
LEOS	low Earth-orbiting satellite
LER	label edge router
LES	loop emulation service
LIDB	line information database
LL	long line
LLC	logical link control
LMDS	local multipoint distribution system
LMN	local network management
LMOS	loop maintenance operation system
LMP	link management protocol
LMS	loop-management system OR loop-monitoring system OR link-monitoring system
LNNI	LANE network-to-network interface
LNP	local number portability
LNS	L2TP network server
LOI	letter of intent
LOL	loss of lock
LOS	line of sight OR loss of signal
LPF	low-pass filter
LQ	listening quality
LRN	local routing number
LSA	label switch assignment OR link state advertisement
LSB	location-sensitive billing
LSMS	local service management system
LSO	local service office
LSP	label-switched path
LSR	label-switched router OR leaf setup request OR local service request
LT	line terminator OR logical terminal
LTE	lite terminating equipment
LUNI	LANE user network interface
LX	local exchange
M2PA	message transfer protocol 2 peer-to-peer adaptation
M2UA	message transfer protocol 2–user adaptation layer
M3UA	message transfer protocol 3–user adaptation layer
MAC	media access control
MADU	multiwave add/drop unit
MAN	metropolitan-area network
MAP	mobile applications part
MAS	multiple-application selection
MB	megabyte
Mb	megabit
MBAC	measurement-based admission control
MBGP	multicast border gateway protocol
MBone	multicast backbone
Mbps	megabits per second
MC	multipoint controller
MCC	mobile country code
MCU	multipoint control unit
MDF	main distribution frame
MDSL	multiple DSL

MDTP	media device transport protocol
MDU	multiple-dwelling unit
MEGACO	media gateway control
MEMS	micro-electromechanical system
MExE	mobile execution environment
MFJ	modified final judgment
MG	media gateway
MGC	media gateway controller
MGCP	media gateway control protocol
MHz	megahertz
MIB	management information base
MII	media independent interface
MIME	multipurpose Internet mail extensions
MIN	mobile identification number
MIPS	millions of instructions per second
MIS	management information system
MITI	Ministry of International Trade and Industry (in Japan)
MLT	mechanized loop testing
MM	Mobility Management
MMDS	multichannel multipoint distribution system
MMPP	Markov-Modulated Poisson Process
MMUSIC	Multiparty Multimedia Session Control [working group]
MNC	mobile network code
MOM	message-oriented middleware
MON	metropolitan optical network
MOP	method of procedure
MOS	mean opinion score
MOSFP	multicast open shortest path first
MOU	minutes of use OR memorandum of understanding
MPC	mobile positioning center
MPEG	Moving Pictures Experts Group
MPI	message passing interface
MPLambdaS	multiprotocol lambda switching
MPLS	multiprotocol label switching
MPOA	multiprotocol over ATM
MPOE	multiple point of entry
MPOP	metropolitan point of presence
MPP	massively parallel processor
MPx	MPEG–Layer x
MRC	monthly recurring charge
MRS	menu routing system
MRSP	mobile radio service provider
ms	millisecond
MSC	mobile switching center
MSF	Multiservice Switch Forum
MSIN	mobile station identification number
MSNAP	multiple services network access point
MSO	multiple-system operator
MSP	management service provider
MSS	multiple-services switching system
MSSP	mobile satellite service provider
MTA	message transfer agent
MTBF	mean time between failures
MTP–[x]	message transfer part–[x]
MTTR	mean time to repair
MTU	multiple-tenant unit
MVL	multiple virtual line
MWIF	Mobile Wireless Internet Forum
MZI	Mach-Zender Interferometer
NA	network adapter
NAFTA	North America Free Trade Agreement
NANC	North American Numbering Council
NANP	North American Numbering Plan
NAP	network access point
NARUC	National Association of Regulatory Utility Commissioners
NAS	network access server
NASA	National Aeronautics and Space Administration
NAT	network address translation
NATA	North American Telecommunications Association
NBN	node-based network
NCF	National Communications Forum
NCP	network control protocol
NCS	national communications system OR network connected server
NDA	national directory assistance
NDM–U	network data management–usage
NDSF	non-dispersion-shifted fiber
NE	network element
NEBS	network-equipment building standards
NEL	network-element layer
NEXT	near-end crosstalk
NFS	network file system
NG	next generation
NGCN	next-generation converged network
NGDLC	next-generation digital loop carrier
NGF	next-generation fiber
NGN	next-generation network
NGOSS	next-generation operations system and software OR next-generation OSS
NHRP	next-hop resolution protocol
NI	network interface
NIC	network interface card
NID	network interface device
NIDS	network intrusion detection system
NIIF	Network Interconnection Interoperability Forum
NIS	network information service
NIU	network interface unit
nm	nanometer
NML	network-management layer
NMS	network-management system
NND	name and number delivery
NNI	network-to-network interface
NOC	network operations center
NOMAD	national ownership, mobile access, and disaster communications
NP	number portability
NPA	numbering plan area
NPAC	Number Portability Administration Center
NPN	new public network
NP–REQ	number-portable request query
NPV	net present value
NRC	Network Reliability Council OR nonrecurring charge
NRIC	Network Reliability and Interoperability Council
NRRI	National Regulatory Research Institute (The Ohio State University)
NRSC	Network Reliability Steering Committee
NRZ	non–return to zero
NS/EP	national security and emergency preparedness

NSAP	network service access point
NSAPI	Netscape server application programming interface
NSCC	network surveillance and control center
NSDB	network and services database
NSP	network service provider OR network and service performance
NSTAC	National Security Telecommunications Advisory Committee
NT	network termination OR new technology
NTN	network terminal number
NTSC	National Television Standards Committee
NVP	network voice protocol
NZ–DSF	nonzero dispersion-shifted fiber
O&M	operations and maintenance
OADM	optical add/drop multiplexer
OA&M	operations, administration, and maintenance
OAM&P	operations, administration, maintenance, and provisioning
OBF	Ordering and Billing Forum
OBLSR	optical bidirectional line-switched ring
OC–[x]	optical carrier–[level x]
OCBT	ordered core-based protocol
OCD	optical concentration device
OCh	optical channel
OCR	optical character recognition
OCU	office channel unit
OCX	open compact exchange
OD	origin-destination
ODBC	open database connectivity
ODSI	optical domain services interface
O–E	optical-to-electrical
O–EC	optical–electrical converter
OECD	Organization for Economic Cooperation and Development
OEM	original equipment manufacturer
O–E–O	optical-to-electrical-to-optical
OEXC	opto-electrical cross-connect
OFC	Optical Fiber Conference
OFDM	orthogonal frequency division multiplexing
OIF	Optical Internetworking Forum
OLA	optical line amplifier
OLAP	on-line analytical processing
OLI	optical link interface
OLT	optical line termination OR optical line terminal
OLTP	on-line transaction processing
OMC	Operations and Maintenance Center
OMG	Object Management Group
OMS SW	optical multiplex section switch
OMS	optical multiplex section
OMSSPRING	optical multiplex section shared protection ring
ONA	open network architecture
ONE	optical network element
ONMS	optical network-management system
ONI	optical network interface
ONT	optical network termination
ONTAS	optical network test access system
ONU	optical network unit
OP	optical path
OPEX	operating expenses
OPS	operator provisioning station
OPTIS	overlapped PAM transmission with interlocking spectra
OPXC	optical path cross-connect
ORB	object request broker
ORT	operational readiness test
OS	operating system
OSA	open service architecture
OSC	optical supervisory panel
OSD	on-screen display
OSGI	open services gateway initiative
OSI	open systems interconnection
OSMINE	operations systems modification of intelligent network elements
OSN	optical-service network
OSNR	optical signal-to-noise ratio
OSP	outside plant
OSPF	open shortest path first
OSS	operations support system
OSS/J	OSS through Java
OSU	optical subscriber unit
OTM	optical terminal multiplexer
OTN	optical transport network
OUI	optical user interface
O-UNI	optical user-to-network interface
OUSP	optical utility services platform
OVPN	optical virtual petabits network OR optical virtual private network
OWSR	optical wavelength switching router
OXC	optical cross-connect
PABX	private automatic branch exchange
PACA	priority access channel assignment
PACS	picture archiving communications system
PAL	phase alternate line
PAM	Presence and Availability Management [Forum] OR pulse amplitude modulation
PAMS	perceptual analysis measurement system
PAN	personal access network
PANS	pretty amazing new services
P&L	profit and loss
PBN	point-to-point–based network OR policy-based networking
PBX	private branch exchange
PC	personal computer
PCF	physical control field
PCI	peripheral component interconnect
PCM	pulse code modulation
PCN	personal communications network
PCR	peak cell rate
PCS	personal communications service
PDA	personal digital assistant
PDD	post-dial delay
PDE	position determination equipment
PDH	plesiochronous digital hierarchy
PDN	public data network
PDP	policy decision point
PDSN	packet data serving node
PDU	protocol data unit
PE	provider edge
PERL	practical extraction and report language
PESQ	perceptual evolution of speech quality
PFD	phase-frequency detector
PHB	per-hop behavior
PHP	personal home page
PHY	physical layer

PIC	point-in-call OR predesignated interexchange carrier OR primary interexchange carrier
PICS	plug-in inventory control system
PIM	personal information manager OR protocol-independent multicast
PIN	personal identification number
PINT	PSTN and Internet Networking [IETF working group]
PINTG	PINT gateway
PKI	public key infrastructure
PLA	performance-level agreement
PLC	planar lightwave circuit OR product life cycle
PLCP	physical layer convergence protocol
PLL	phase locked loop
PLMN	public land mobile network
PLOA	protocol layers over ATM
PM	performance monitoring
PMD	physical-medium dependent OR polarization mode dispersion
PMDC	polarization mode dispersion compensator
PMO	present method of operation
PMP	point-to-multipoint
PN	personal number
PNNI	private network-to-network interface
PnP	plug and play
PO	purchase order
PODP	public office dialing plan
POET	partially overlapped echo-cancelled transmission
POF	plastic optic fiber
POH	path overhead
POIS	packet optical interworking system
PON	passive optical network
POP	point of presence
POP3	post office protocol 3
POS	packet over SONET OR point of service
PosReq	position request
POT	point of termination
POTS	plain old telephone service
PP	point-to-point
PPD	partial packet discard
PPP	point-to-point protocol
PPPoA	point-to-point protocol over ATM
PPPoE	point-to-point protocol over Ethernet
PPTP	point-to-point tunneling protocol
PP–WDM	point-to-point–wavelength division multiplexing
PQ	priority queuing
PRI	primary rate interface
ps	picosecond
PSAP	public safety answering point
PSC	Public Service Commission
PSD	power spectral density
PSDN	public switched data network
PSID	private system identifier
PSN	public switched network
PSPDN	packet-switched public data network
PSQM	perceptual speech quality measure
PSTN	public switched telephone network
PTE	path terminating equipment
PTN	personal telecommunications number service
PTP	point-to-point
PTT	Post Telephone and Telegraph Administration
PUC	public utility commission
PVC	permanent virtual circuit
PVM	parallel virtual machine
PVN	private virtual network
PWS	planning workstation
PXC	photonic cross-connect
QAM	quadrature amplitude modulation
QOE	quality of experience
QoS	quality of service
QPSK	quaternary phase shift keying
QSDG	QoS Development Group
RADIUS	remote authentication dial-in user service
RADSL	rate-adaptive DSL
RAM	remote access multiplexer
RAN	regional-area network
RAP	resource allocation protocol
RAS	remote access server
RBOC	regional Bell operating company
RCP	remote call procedure
RCU	remote control unit
RDBMS	relational database management system
RDC	regional distribution center
RDSLAM	remote DSLAM
REL	release
RF	radio frequency
RFC	request for comment
RFI	request for information
RFP	request for proposal
RFPON	radio frequency optical network
RFQ	request for quotation
RGU	revenue-generating unit
RGW	residential gateway
RHC	regional holding company
RIAC	remote instrumentation and control
RIM	research in motion
RIP	routing information protocol
RISC	reduced instruction set computing
RJ	registered jack
RLL	radio in the loop
RM	resource management
RMA	request for manual assistance
RMI	remote method invocation
RMON	remote monitoring
ROADM	reconfigurable optical add/drop multiplexer
ROBO	remote office/branch office
ROI	return on investment
ROW	right of way
RPC	remote procedure call
RPF	reverse path forwarding
RPR	resilient packet ring
RPRA	Resilient Packet Ring Alliance
RPT	resilient packet transport
RQMS	requirements and quality measurement system
RRQ	round-robin queuing or registration request
RSU	remote service unit
RSVP	resource reservation protocol
RSVP–TE	resource reservation protocol–traffic engineering
RT	remote terminal

RTCP	real-time conferencing protocol
RTOS	real-time operating system
RTP	real-time transport protocol
RTSP	real-time streaming protocol
RTU	remote test unit
RxTx	receiver/transmitter
RZ	return to zero
SAM	service access multiplexer
SAN	storage-area network
SAP	service access point
SAR	segmentation and reassembly
S-band	short band
SBS	stimulated Brillouin scattering
SCAN	switched-circuit automatic network
SCCP	signaling connection control part
SCCS	switching control center system
SCE	service-creation environment
SCF	service control function
SCL	service control language
SCM	service combination manager OR station class mark OR subscriber carrier mark
SCN	service circuit node OR switched-circuit network
SCP	service control point
SCR	sustainable cell rate
SCSP	server cache synchronization protocol
SCTP	simple computer telephony protocol OR simple control transport protocol OR stream control transmission protocol
SD	selective discard
SDA	separate data affiliate
SD&O	service development and operations
SDB	service design bureau
SDC	service design center
SDF	service data function
SDH	synchronous digital hierarchy
SDM	service-delivery management OR shared data model
SDN	software-defined network
SDP	session description protocol
SDRP	source demand routing protocol
SDSL	symmetric DSL
SDTV	synchronous digital hierarchy
SDV	switched digital video
SE	service element
SEC	Securities and Exchange Commission
SEE	service-execution environment
SEP	signaling endpoint
ServReq	service request
SET	secure electronic transaction
SFA	sales force automation
SFD	start frame delimiter
SFF	small form-factor
SFGF	supplier-funded generic element
SG	signaling gateway
SG&A	selling, goods, and administration OR sales, goods, and administration
SGCP	simple gateway control protocol
SGSN	serving GPRS support node
SHLR	standalone home location register
SHV	shareholder value
SI	systems integrator
SIBB	service-independent building block
SIC	service initiation charge
SICL	standard interface control library
SID	silence indicator description
SIF	SONET Interoperability Forum
sigtran	Signaling Transport [working group]
SIM	subscriber identity module OR service interaction manager
SIP CPL	SIP call processing language
SIP	session initiation protocol
SIP–T	session initiation protocol for telephony
SITA	Société Internationale de Télécommunications Aéronautiques
SIU	service interface unit
SIVR	speaker-independent voice recognition
SKU	stock-keeping unit
SL	service logic
SLA	service-level agreement
SLC	subscriber line carrier
SLEE	service logic execution environment
SLIC	subscriber line interface circuit
SLO	service-level objective
SM	sparse mode
SMC	service management center
SMDI	simplified message desk interface
SMDS	switched multimegabit data service
SME	small-to-medium enterprise
SMF	single-mode fiber
SML	service management layer
SMP	service management point
SMPP	short message peer-to-peer protocol
SMS	service-management system OR short message service
SMSC	short messaging service center
SMTP	simple mail transfer protocol
SN	service node
SNA	service node architecture OR service network architecture
SNAP	subnetwork access protocol
SNMP	simple network-management protocol
SNPP	simple network paging protocol
SNR	signal-to-noise ratio
SO	service objective
SOA	service order activation
SOAC	service order analysis and control
SOAP	simple object access protocol
SOCC	satellite operations control center
SOE	standard operating environment
SOHO	small office/home office
SON	service order number
SONET	synchronous optical network
SOP	service order processor
SP	service provider OR signaling point
SPC	stored program control
SPE	synchronous payload envelope
SPF	shortest path first
SPIRITS	Service in the PSTN/IN Requesting Internet Service [working group]
SPIRITSG	SPIRITS gateway
SPM	self-phase modulation OR subscriber private meter
SPOP	service point of presence
SPX	sequence packet exchange
SQL	structured query language
SRF	special resource function
SRP	source routing protocol
SRS	stimulated Raman scattering
srTCM	single-rate tri-color marker

SS	softswitch
SS7	signaling system 7
SSE	service subscriber element
SSF	service switching function
SSG	service selection gateway
SSL	secure sockets layer
SSM	service and sales management
SSMF	standard single-mode fiber
SSP	service switching point
STE	section terminating equipment
STM	synchronous transfer mode
STN	service transport node
STP	shielded twisted pair OR signal transfer point OR spanning tree protocol
STR	signal-to-resource
STS	synchronous transport signal
SUA	SCCP user adaptation
SVC	switched virtual circuit
SW	software
SWAN	storage wide-area network
SWAP	shared wireless access protocol
SWBT	Southwestern Bell Telephone
SWOT	strengths, weaknesses, opportunities, and threats
SYN	IN synchronous transmission
TALI	transport adapter layer interface
TAPI	telephony application programming interface
TAT	terminating access trigger OR termination attempt trigger OR transatlantic telephone cable
Tb	terabit
TBD	to be determined
Tbps	terabits per second
TC	tandem connect
TCAP	transactional capabilities application part
TCB	transfer control block
TCIF	Telecommunications Industry Forum
TCM	time compression multiplexing
TCO	total cost of ownership
TCP	transmission control protocol
TCP/IP	transmission control protocol/Internet protocol
TC–PAM	trellis coded–pulse amplitude modulation
TDD	time division duplex
TDM	time division multiplex
TDMA	time division multiple access
TDMDSL	time division multiplex digital subscriber line
TDOA	time difference of arrival
TDR	time domain reflectometer OR transaction detail record
TE	traffic engineering
TEAM	transport element activation manager
TED	traffic engineering database
TEM	telecommunications equipment manufacturer
TFD	toll-free dialing
THz	terahertz
TIA	Telecommunications Industry Association
TIMS	transmission impairment measurement set
TINA	Telecommunications Information Networking Architecture
TINA-C	Telecommunications Information Networking Architecture Consortium
TIPHON	Telecommunications and Internet Protocol Harmonization over Networks
TIWF	trunk interworking function
TL1	transaction language 1
TLDN	temporary local directory number
TLS	transparent LAN service OR transport-layer security
TLV	tag length value
TMF	TeleManagement Forum
TMN	telecommunications management network
TMO	trans-metro optical
TN	telephone number
TNO	telecommunications network operator
TO&E	table of organization and equipment
TOD	time of day
TOM	telecom operations map
ToS	type of service
TP	twisted pair
TPM	transaction processing monitor
TPS–TC	transmission control specific–transmission convergence
TR	technical requirement OR tip and ring
TRA	technology readiness assessment
TRIP	telephony routing over Internet protocol
trTCM	two-rate tri-color marker
TSB	telecommunication system bulletin
TSI	time slot interchange
TSP	telecommunications service provider
TSS	Telecommunications Standardization Section
TTC	Telecommunications Technology Committee
TTCP	test TCP
TTL	transistor-transistor logic
TTS	text-to-speech OR TIRKS® table system
TUI	telephone user interface
TUP	telephone user part
TV	television
UA	user agent
UADSL	universal ADSL
UAK	user-authentication key
UAWG	Universal ADSL Working Group
UBR	unspecified bit rate
UBT	ubiquitous bus technology
UCP	universal computer protocol
UCS	uniform communication standard
UDP	user datagram protocol
UI	user interface
ULH	ultra-long-haul
UM	unified messaging
UML	unified modeling language
UMTS	Universal Mobile Telecommunications System
UN	United Nations
UNE	unbundled network element
UNH–IOL	University of New Hampshire–Interoperability Laboratory
UNI	user network interface
UOL	unbundled optical loop
UPC	usage parameter control
UPI	user personal identification
UPS	uninterruptible power supply
UPSR	unidirectional path-switched ring
URI	uniform resource identifier
URL	universal resource locator

USB	universal serial bus
USTA	United States Telephone Association
UTOPIA	Universal Test and Operations Interface for ATM
UTS	universal telephone service
V&H	vertical and horizontal
VAD	voice activity detection
VAN	value-added network
VAR	value-added reseller
VAS	value-added service
VASP	value-added service provider
VBNS	very–high-speed backbone network service
VBR	variable bit rate
VBR–nrt	variable bit rate–non–real-time
VBR–rt	variable bit rate–real time
VC	virtual circuit OR virtual channel
VCC	virtual channel connection
VCI	virtual channel identifier
VCLEC	voice CLEC
VCO	voltage-controlled oscillator
VCR	videocassette recorder
VCSEL	vertical cavity surface emitting laser
VD	visited domain
VDM	value delivery model
VDSL	very-high–data rate DSL
VeDSL	voice-enabled DSL
VGW	voice gateway
VHE	virtual home environment
VHS	video home system
VITA	virtual integrated transport and access
VLAN	virtual local-area network OR voice local-area network
VLR	visitor location register
VLSI	very–large-scale integrated
VM	virtual machine
VMS	voice-mail system
VoADSL	voice over ADSL
VoATM	voice over ATM
VOD	video on demand
VoDSL	voice over DSL
VoFR	voice over frame relay
VoIP	voice over IP
VON	voice on the Net
VoP	voice over packet
VOQ	virtual output queuing
VoT1	voice over T1
VP	virtual path
VPDN	virtual private dial network
VPI	virtual path identifier
VPIM	voice protocol for Internet messaging
VPN	virtual private network
VPR	virtual path ring
VPRN	virtual private routed network
VRE	Virtual Radiology Environment
VRU	voice response unit
VSAT	very-small–aperture terminal
VSI	virtual switch interface
VSM	virtual services management
VSN	virtual service network
VSR	very short reach
VT	virtual tributary
VTN	virtual transport network
VToA	voice traffic over ATM
VVPN	voice virtual private network
VXML	voice extensible markup language
W3C	World Wide Web Consortium
WAN	wide-area network
WAP	wireless application protocol
WATS	wide-area telecommunications service
WB DCS	wideband DCS
WCDMA	wideband CDMA
WCT	wavelength converting transponder
WDCS	wideband digital cross-connect
WDM	wavelength division multiplexing
WECA	wireless Ethernet compatibility alliance
WEP	wired equivalent privacy
WFA	work and force administration
WFQ	weighted fair queuing
WIM	wireless instant messaging
WIN	wireless intelligent network
WLAN	wireless local-area network
WLL	wireless local loop
WMAP	wireless messaging application programming interface
WML	wireless markup language
WNP	wireless local number portability
WRED	weighted random early discard
WS	work station
WSP	wireless session protocol
WTA	wireless telephony application
WUI	Web user interface
WVPN	wireless VPN
WWCUG	wireless/wireline closed user group
WWW	World Wide Web
XA	transaction management protocol
XC	cross-connect
XD	extended distance
xDSL	[see DSL]
XML	extensible markup language
XNS	Xerox network system
XPM	cross-phase modulation
XPS	cross-point switch
xSP	specialized service provider
XT	crosstalk
XTP	express transport protocol
Y2K	Year 2000

Index