◆ Next-Generation Networks

Douglas C. Dowden, Richard D. Gitlin, and Robert L. Martin

Dramatic changes in technology, network services, and regulation are driving the rapid definition and implementation of new network architectures and elements. This paper discusses the forces that are moving today's networks toward converged, next-generation networks. It also provides a multiservice, high-performance, and highly reliable next-generation network architecture. This architecture serves as a framework for placing the papers in this issue in context. Together, these papers provide a view of Lucent Technologies' Network Vision with a focus on packet networking, including network architecture and planning, intelligent packet switching, packet voice, quality of service, virtual private networks, and network and service management.

Introduction

Dramatic changes in technology, network services, and regulation are driving the rapid definition and implementation of new network architectures and elements. The inexorable march of technology in microelectronics, software, photonics, and wireless is disrupting the fundamental nature of networks and convulsing the markets they serve. Concurrently, the dramatically increasing demand for high-speed, high-capacity data services, due in large part to the popularity of the Internet and the ubiquity of the IP network-level protocol with its rapidly maturing associated protocols and standards, is driving the next-generation network architecture toward a packet network. The consequences of this shift are far greater than the conversion from an analog to a digital network that was the focus of network reengineering in recent decades. The architecture and distributed control of next-generation packet networks provide enormous efficiencies via statistical multiplexing of user traffic, ensure that end users can readily upgrade and introduce new applications, provide advanced services by decoupling of the user and system transmission rates, and promise simplified and unified management systems and interfaces. Next-generation networks provide a platform for rapid service creation, by incorporating intelligence in separate standards-based network servers that build on rapid technological advances in service-enabling software, software-controlled switching, client/server computing, and new database/directory systems.

While many of the network directions are quite predictable, there is still significant uncertainty in the markets that will ultimately be served and in fundamental approaches necessary to assure carrier-grade service for large-scale networks and advanced network services. The major business challenges facing next-generation network providers and suppliers are:

- To build networks that provide levels of quality and reliability that meet or exceed those in the voice network,
- To match customer expectations for networks that can support the rapid introduction of a rich new suite of services, and
- To create networks that are cost effective and easy to manage.

Given these challenges, next-generation networks will require a high degree of agility in architectural approach and software programmability. For the circuit network, an intelligent network has been developed, and it continues to evolve, allowing more services to be offered. A similar, but more powerful, Internet protocol (IP) services platform will evolve for

the next-generation network. This platform will build on service level agreements (SLAs) between customers and service suppliers, and the network will execute these agreements through policies that tell the devices and their support systems how to process, route, and switch data packets. The management of these policies is the key to providing reliable, scalable, secure, and manageable data networks. Significant business and technical challenges exist in making the transition from today's networks to next-generation, packet-based networks.

Forces of Change

Several compelling events and trends—network convergence, disruptive technologies, and convulsive markets—have already begun to drive today's networks toward next-generation networks.

Network Convergence and Collaboration

Two major networks of networks, the public switched telephone network (PSTN) and the Internet, exist today. Each is tuned to the anchor services it supports. The PSTN is based on circuit switches that provide a very high quality of service for large-scale, advanced voice services—for example, free-phone, operator services, and customer calling. Data services—for example, fax and e-mail—operate over this network utilizing circuit connections. The PSTN is a low-delay, fixed-bandwidth network of networks optimally matched to the mathematically modeled, well-behaved Poisson voice traffic it supports. PSTN services are provided from switches, adjuncts to switches, and the intelligent network. Wireless networks—which mainly provide mobility voice services today—interconnect to the PSTN, providing seamless global connectivity for all wired and wireless users.

The Internet is based mainly on packet switches that provide very flexible data services—for example, e-mail, virtual private networks, and access to the World Wide Web. At the same time, however, it is a variable-delay, variable-bandwidth network that provides no guarantees on the quality of service (QoS) that the packet-based traffic it supports will experience. The explosive growth of the Internet and related intranets (private packet networks operated by organizations) and extranets (networks shared among sev-

---

Panel 1. Abbreviations, Acronyms and Terms

ADSL—asymmetric digital subscriber line
ASIC—application specific integrated circuit
ATM—asynchronous transfer mode
DiffServ—differentiated services
DWDM—dense wave division multiplexed/multiplexing
H.323—suite of protocols for audio/video communication over non-QoS networks such as IP
ILS—Internet location server
IP—Internet protocol
IPX—internetwork packet exchange
ITU—International Telecommunication Union
LAN—local area network
LDAP—lightweight directory access protocol
MGCP—media gateway control protocol
MPLS—multiprotocol label switching
PSTN—public switched telephone network
QoS—quality of service
RAM—random access memory
RAS—registration/administration/status protocol
RSVP—resource reservation protocol
SDH—synchronous digital hierarchy
SIP—session initiation protocol
SLA—service level agreement
SNA—systems network architecture
SONET—synchronous optical network
SS7—Signaling System 7
STM—synchronous transfer mode
TINA-C—Telecommunications Information Networking Architecture Consortium
TMN—Telecommunications Management Network
VPN—virtual private network

---

eral groups such as suppliers and customers) is such that it is projected that data traffic volumes will pass voice traffic volumes around the world over the next few years. It is also projected that the number of devices connected to the Internet will equal the number of people on the globe in eight to ten years. These projections become fathomable if we consider that in the next few years many electronic appliances will have IP addresses, creating enormous potential for machine-to-machine and person-to-machine data communications and collaboration.

As data traffic surpasses voice traffic, and as it becomes possible to provide high levels of QoS on

packet networks—particularly for voice and other real-time services—it will be desirable to converge the multiple networks around a single packet-based core network. Such convergence will support emerging multimedia services, increase the ability of a carrier to support the multiple needs of its customers, and reduce the cost of network operations. This convergence around a packet-based core will also allow the many networks to collaborate so that customers will perceive they are working with a single, integrated network. Such convergence will likely repeat itself as enterprise and public networks similarly converge on a standard set of protocols, policies, and architectures.

Disruptive Technologies

Technology advances can be described by learning curves, which plot a capacity characteristic of a technology over time. When the advance is dramatic and sustained, the technology can have a disruptive impact on the markets it creates and serves.[1] Continuing exponentially increasing advances in microelectronics, photonics, and wireless technology are supporting the market drive to converged networks. These technologies provide the cost-effective foundation for the network elements that will make up the next-generation converged networks. It will be software, new architectures and algorithms, and QoS-based designs, however, that will make these networks carrier grade, and thus desirable.

Semiconductor technology continues to follow the famous Moore's Curve whereby the number of transistors on a chip doubles every 18 months. This doubling, which should continue until roughly the year 2010, has been the technology fuel for improving the price performance of the computer industry. A major impact of this trend will be "systems on chip," a design paradigm that will lead to continuing functional integration and reuse opportunities in network elements, similar to the design paradigm that is taking place in the PC industry. The use of very high-density custom silicon chips will be critical in realizing the advanced algorithms needed for large-scale, high-speed, QoS-aware network elements.

Optical transmission capacity is doubling every 12 months through increasing the number of wave lengths of light on a single fiber—dense wavelength division multiplexing—and increasing the number of bits carried by a single wavelength. Experimental systems can carry the equivalent of all of the voice traffic in the world on a single, long-distance fiber—that is, 1 terabit for 400 kilometers. The underlying technology for this single-fiber experiment easily scales to 10 terabits while the underlying physics supports the order of 100 terabits. Such huge capacity has significant implications on network architectures. It will be possible to build highly distributed, advanced network services. Intelligence will appear in many locations—in devices and network appliances connected to the network, in servers at the edge of the network, and in servers deep in the network.

Wireless capacity in a given volume of air is doubling every nine months through advanced signal processing and intelligent antenna and receiver technology. Such technology advances will lead to economically feasible fixed wireless loops. These loops will be among the leading economic choices for narrow bandwidth access in developing countries and, over time, similarly for new broadband access in developed countries. This choice will be driven by the decreasing cost of wireless systems due to the technology curves, the low labor and first costs to install them, and the ease and selectivity of their deployment.

The above technological breakthroughs have produced processors and memories that are consonant with advanced networking. Specifically, the rate of improvement in processor speeds, random access memory (RAM) capacity, and memory access speeds have all grown at the same rate as the fiber-optic speeds available to the end user. This implies that the sophisticated control processing and memories will exist to support the rapid bursts of arriving high-speed packets in network elements.

Together, these technologies will make it possible to build networks having roughly 100 to 250 times the capacity of today's networks at the cost of today's networks by the year 2005—as long as the scaling and QoS problems already mentioned can be addressed.

Convulsive Markets

As the computer industry has taught us, markets are far less predictable than technology. The forces of change are far more complex than technology itself,

and the impact of innovation in appealing applications is impossible to forecast. Two forces are clear, however. First, telecommunications markets around the world are being deregulated, privatized, and opened for competition. New carriers are entering all markets—wired, wireless, and data. Second, these new carriers' networks are increasingly targeted toward converged networks that attack the markets of the embedded carriers. Likewise, incumbent service providers are building very competitive, advanced networks that build on their core capabilities. The market battle for the converged networks will be both mammoth and complex.

Next-Generation Network Architecture

As mentioned earlier, different types of networks exist today. It is likely, however, that next-generation networks will share a high-level architecture, as shown in **Figure 1**. Differing in detail from carrier to carrier, migration to this architecture will likely be the most complex aspect of realizing future converged networks. While Moore's Law and the dramatic advances in other fundamental technologies provide the underpinnings for the emerging new networks, it has been recognized—in what is sometimes called Metcalfe's Law—that the value of a network grows exponentially with the number of users and interconnected sources. Fueled by the lower cost and the increasing availability of bandwidth, data networks will provide increased levels of connectivity and performance. These networks will be arranged as collaborating networks of networks that become the organizing principle for most communications. They will contain the following major ingredients:

- A shared, packet-based, optical core network using DWDM optical transport with optical add/drop and multiplexing.
- A variety of broadband access mechanisms—wired and wireless—for all types of customers. Higher-speed access will eliminate the access bottleneck for advanced services.
- Access to traditional and new advanced network services provided by interworking to the PSTN and to new open applications platforms.
- A new style of network management based on active directories and policy managers.

With these attributes, the new networks will be able to support both the demanding applications that are emerging today—such as virtual private networks (VPN's), packet voice, collaboration, electronic commerce, and remote data access—and the innovative applications of tomorrow.

The remainder of this section expands on these ingredients.

Packet-Based Optical Core

The core transport network will be built around DWDM transport systems. Systems carrying 400Gb/s utilizing 80 wavelengths are now available. They will rapidly evolve to carry significantly more information using an ever-increasing number of wavelengths.

Multiterabit offices, based on multiprotocol (for example, asynchronous tranfer mode [ATM] and Internet protocol [IP]) bandwidth management systems and multiplexers, will be connected via these DWDM systems. The optical core will evolve over time to utilize optical add/drop multiplexers and cross connects. This "slow" optical switching will reduce the cost of networks by reducing the number of optical/electrical conversions and will migrate bandwidth management toward wavelength management.

Bandwidth management systems illustrate both the power of advanced microelectronics and the realization of an important design approach—architectural agility. A bandwidth management system is an integration of traditional cross-connect systems and multiplexers that support synchronous transfer mode (STM)—as SONET or synchronous digital hierarchy (SDH)—ATM, or IP directly. The ability to support multiple protocols allows a carrier to build a network based on the technical characteristics and market appeal of the protocols as they exist today. More importantly, it allows the carrier to migrate the network as the protocols evolve—in both technical capability and market acceptance.

Protocol evolution, particularly as it addresses the relative roles of ATM and IP, is a hotly debated topic. Circuit-switched networks are mainly tuned to the range of today's voice-band applications. The inherent flexibility in packet networks that allows support of multiple applications also makes it substantially more
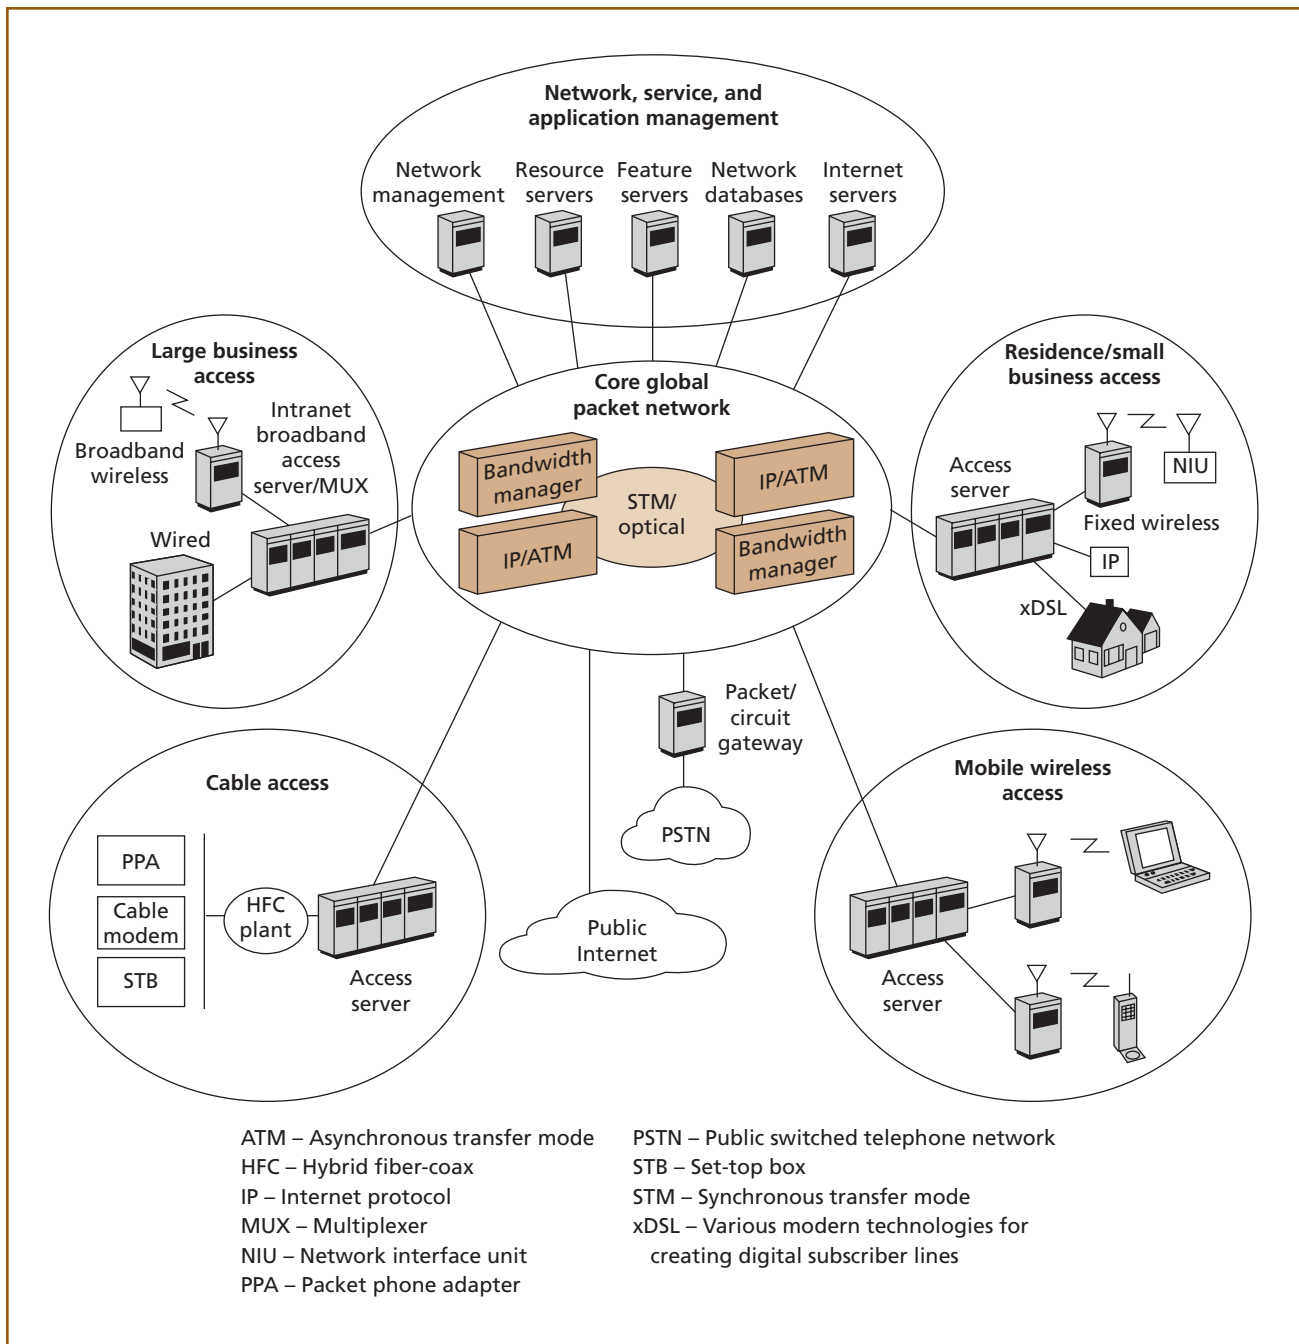
Network, service, and application management

Network management Resource servers Feature servers Network databases Internet servers

Core global packet network

Large business access

Broadband wireless

Intranet broadband access server/MUX

Wired

Bandwidth manager

STM/optical

IP/ATM

IP/ATM

Bandwidth manager

Residence/small business access

Access server

NIU

Fixed wireless

IP

xDSL

Packet/circuit gateway

PSTN

Cable access

PPA

Cable modem

STB

HFC plant

Access server

Public Internet

Mobile wireless access

Access server

Access server

ATM – Asynchronous transfer mode
HFC – Hybrid fiber-coax
IP – Internet protocol
MUX – Multiplexer
NIU – Network interface unit
PPA – Packet phone adapter

PSTN – Public switched telephone network
STB – Set-top box
STM – Synchronous transfer mode
xDSL – Various modern technologies for
creating digital subscriber lines

Figure 1.
The real world of communications networking.

difficult to provide QoS to a large set of users and applications. Today, building a converged network core that provides QoS requires accommodation of four protocols—IP, STM, ATM, and DWDM. Each one provides important network capabilities: IP provides data interoperability and emerging support for QoS, ATM provides guaranteed QoS support for multiple services, STM provides reliable transport, and DWDM provides high-speed transport. It is likely that these four protocols will converge over time to two protocols—one dealing with the lower layer physical domain and the other dealing with network layer issues.

Protocols will converge toward one called "IP" for market appeal. This protocol will capture today's IP and ATM capabilities. Major progress is rapidly being made to provide nearly the same of level of QoS in IP and ATM-based networks. Technology advances to mitigate the inherent delay and jitter problems present in today's IP-based networks are occurring at a rapid rate. These advances include new network design algorithms—for example, in this issue, "VPN DESIGNER: A Tool for Design of Multiservice Private Networks" by Mitra, Morrison, and Ramakrishnan[2] describes a software tool for design of intranets or virtual private networks on a service provider's infrastructure. The advances also include sophisticated edge and core intelligent switches that provide new levels of QoS support (see, for example, "The PacketStar™ 6400 IP Switch—An IP Switch for the Converged Network" by Fendick et al.[3]) and new standards such as multiprotocol label switching (MPLS), resource reservation protocol (RSVP), and differentiated services (DiffServ) that support the required QoS across a wide set of applications and network domains. In "QoS with Differentiated Services," Weiss[4] describes the fundamental components of QoS and how they can be supported through various IP-based solutions. Similarly, a new, simple and cost-effective optical protocol supplanting today's STM and DWDM protocols will likely emerge.

In the interim, while both the technology and the market evolve, it is crucial that a carrier continue toward building the converged core in a way that provides agility for uncertain and unpredictable technology and market evolution. The paper by Doshi et al.—"A Comparison of Next-Generation IP-Centric Transport Architectures"[5]—evaluates transport architecture alternatives to carry IP-based traffic in the backbone.

Packet switches—both ATM and IP—will be connected to the optical core. The interconnection protocols will evolve. DWDM, as well as ATM and STM, will carry IP. ATM switches, which are based on connection-oriented protocols, have evolved from the public network marketplace. Their design focus and market base is targeted toward high QoS services. IP switches, which are based on connectionless-oriented protocols, have evolved from routers and the enterprise market-

place. While the design focus and market base of IP routers has been targeted toward low QoS data services, IP switches are using hardware-realized, novel packet classifying, scheduling, and queuing algorithms to serve the high QoS market space.

Though quite different, these two families of packet switches share several key architectural aspects. First, each family consists of scalable product lines utilizing common software for network management, protocol support, and advanced network services support. Second, each family's price performance is rapidly increasing via the use of high-speed microprocessors and custom ASIC's. Third, each family has "edge" and "core" family members. The edge members, based either on IP or ATM, provide protocol conversion for the legacy protocols existing in today's networks—for example SNA, frame relay, or IPX. This legacy support allows transition from existing to next-generation networks. The core switches provide the capacity and, today, are built around a single protocol. With advancing silicon technology, multi-protocol switches will emerge supporting both ATM and IP. Such technology will provide the market/technology agility for packet switching in the same way as the bandwidth manager provides agility for backbone multiplexing and transport.

Broadband Access

There will be multiple broadband access solutions for residence and business. This is in contrast to the core of the network, where a "single solution" is emerging. A particular access solution will be selected by carriers to optimize their existing access systems, if any, and to best match their proposed network services. Access systems provide three capabilities—high-speed transport, multiplexing/routing, and embedded network migration. Even more than in switching, it is crucial that protocol agility be provided in access systems. The access system is where the customer's choice of fundamental protocol—ATM or IP—is first and, perhaps, most importantly seen. The access system alternatives are described below.

ADSL protocols allow a few megabits to be carried over the existing copper plant cost effectively. These protocols support both traditional voice services and data-oriented Internet connections. Digital signal pro-

cessing technology is used to implement algorithms for rapid, cost-effective transport over much of the existing plant. Advanced access systems multiplex and route data traffic to packet networks and voice traffic to either the PSTN, integrated services digital network (ISDN), or a packet network. This multiplexing/routing capability can provide the anchor point around which access systems will evolve. In "The PathStar™ Access Server: Facilitating Carrier-Scale IP Telephony," Fossaceca, Sandoz, and Winterbottom[6] describe an approach to providing a highly cost-effective end-to-end data and voice solution, terminating subscriber loops, providing routing and switching, and offering administrative and billing tools.

In a similar vein, a cable modem allows a few megabits of data to be carried over hybrid fiber-coax cable access systems. The frequency spectrum is split both at the customer site and at the cable head end, thereby allowing access to entertainment, packet data, and, optionally, PSTN services.

Fixed wireless systems, for both residential and business access, are potentially the most disruptive of the access technologies. These systems potentially offer very high, selectively deployable bandwidth. However, they are the most difficult systems to forecast because of the multiple parameters underlying their evolution and use. They are technically complex as they follow silicon and the wireless learning curves. Further, they utilize new propagation technology through intelligent antennas and receivers. Implementation of this new technology requires deeper understanding of propagation characteristics in rural, suburban, and urban settings than now exists. Finally, as in all wireless systems, their use depends on the availability and cost of spectrum as well as the regulatory climate surrounding antenna construction.

Optical access options exist as well. In the near term, the packet-based optical core will repeat itself, albeit with lower capacity, to provide metropolitan transport and business access to the core network described above. In the intermediate term, as the costs for optical/electrical conversion devices drop, these system architectures will repeat themselves yet again to provide optical LANs for business. Likewise, passive optical networks will likely emerge for residential and small business access. These systems offer virtually unlimited bandwidth, and thus are best positioned to support the full range of services that can be envisioned today. The challenge of these systems lies in producing cost-effective optical devices and in determining cost-effective and nonintrusive deployment approaches.

Advanced Network Services

The packet-based optical core and the broadband access systems will provide a low-cost, low-delay broadband infrastructure. A network services architecture must be integrated with this infrastructure to provide seamless interworking to PSTN services and to provide open platforms for new services. The network services architecture must hide the complexity and evolution of the infrastructure. Key elements of this architecture—gateways and gatekeepers, network servers, and network management—are described below.

Gateways and gatekeepers. Gateways and gatekeepers allow packet network customers seamless access to PSTN services, particularly those of the intelligent network. Gateways provide the bearer traffic and signaling conversions between the PSTN and the IP/ATM packet networks. Gatekeepers administer gateways and provide for interworking between domains managed by gatekeepers. This interworking is a crucial ingredient for network migration, convergence, and collaboration.

The subtly complex gateway conversions are the linchpin for flexible and cost-effective interworking. Gateways must support the explosion of signaling standards (for example, H.323, SS7, SIP, and MGCP) and directories (for example, SIP, ILS, RAS, and LDAP) to assure interoperability and collaboration between the many different types of networks being deployed. Gateways must also provide bearer circuit conversions for both interworking and signal compression. "Voice over PacketStar™ Gateway Solution for Service Provider Networks" by Whang et al.[7] discusses issues related to bridging circuit and packet networks to efficiently move voice traffic and voice-related services over packet networks.

Network servers. Two differing approaches—both centered on applications platforms residing on network servers—are emerging for advanced net-

work services. One approach is emerging from the telecommunications industry; the other, from the computer industry. Both will likely continue to develop and will have a strong influence on the converged network architecture.

For several decades, the telecommunications industry has been accelerating network services innovation with a particular architectural approach toward developing and deploying network services. The architectural approach has two ingredients—separation of service management from connection management and open services platforms. The advanced intelligent network, the first instantiation of these architectural principles, has accelerated network services development and deployment. In "Voice Services in Next-Generation Networks: The Evolution of the Intelligent Network and Its Role in Generating New Revenue Opportunities," Kozik, Montgomery, and Stanaway[8] describe three network evolution scenarios where IN-based services play a key role. A thorny technical problem—feature interaction[9]—which can cause service denial and/or confusion, has slowed progress toward the open platform goal. Progress is being made, however, via the Telecommunications Information Networking Architecture Consortium (TINA-C),[10] but much work will be required before it becomes clear whether TINA-C will solve the feature interaction problem.

The Internet and its many servers best exemplify the computer industry approach. Using IP for interoperability, many different servers—each focused on a particular task such as authentication or unified messaging—are being developed and deployed. This approach allows for the rapid introduction of individual services. It does not facilitate the complex interaction of highly related services, but it can be used to package highly related services—for example, toll/tandem services in a feature server. Further, this style of service development and implementation will also be employed directly in the next-generation intelligent terminals—phones, computers, and information appliances.

Finally, Internet-style IP services platforms will emerge, many based on the emerging Internet Engineering Task Force SIP standard. Such platforms will ease the development of applications that build on the interoperability and collaboration of the networks to which they will be connected. These issues are further described in "The Session Initiation Protocol: Providing Advanced Telephony Services Across the Internet" by Schulzrinne and Rosenberg.[11]

Key next-generation network services. Key services to be offered in next-generation networks are driving architecture and design of network elements. Busschbach, in "Toward QoS-Capable Virtual Private Networks,"[12] discusses two methods of implementing QoS-capable VPNs, one over a pure ATM structure and the other based on MPLS, a new protocol that offers a connection-oriented approach to provide QoS support for IP transport over a variety of infrastructures. Providing voice services is a critical issue for next-generation networks. Papers in this issue discuss key issues in implementing voice services. In "Voice over ATM," Chambers et al.[13] describe the economics of this service offering compared to traditional circuit switched networks; and in "Bridging the Gap to IP Telephony," Sijben and Spergel[14] discuss technical, economic, and standards issues connected to voice service offerings on an IP-based network.

Network management. Network Management is an advanced network service. The traditional network management functions—service management, fault isolation, performance management, traffic management, for example—will also be found in next-generation networks. Further, the ITU's Telecommunications Management Network (TMN)[15,16] layered network management architecture will remain a useful model for building open network management systems. However, the traffic characteristics of packet data, the introduction of a variety of network servers, the increasing role of customer premises equipment in network services, and network interworking will place extraordinary new demands on network management. "Network Information Models and OneVision® Architecture" by Lin, Price, and Srinivas[17] describes the collaborative framework of Lucent Technologies' data network management solution.

Next-generation networks will likely use policies to provide the translation from business agreements (for example, SLAs) to the instructions that the net-

works will use to implement services. Policy-based systems will facilitate the offering of advanced applications, such as VPN's and packet voice services. Policy management is largely aimed at providing QoS and security (authentication, authorization, and accounting) using advanced, consolidated directories with a common schema for representing information in conjunction with a policy server. The policy server is responsible for assembling the relevant information, making decisions based on a priori rules and the state of network resources, and then communicating the decisions to the network elements. These policy servers will collaborate with active directories to obtain customer- and network-specific network management criteria. Papers in this issue describe realizations of policy servers to support VPNs and packet voice. Policy management is a key technology to support packet networking across intranets, extranets, VPNs, and possibly the Internet, and standards are likely to develop to support policies across multi-vendor networks.

Technical Challenges

Two technical challenges—QoS and transition—are the cornerstones to realizing viable next-generation networks. It is relatively simple to demonstrate all of the architectural attributes of next-generation networks either in the laboratory or in an enterprise network. The challenge is in scaling them to real public networks that can be deployed and evolved.

Quality of Service (QoS)

Next-generation networks must provide the PSTN's level of QoS. In this context, QoS means availability, reliability, performance, and integrity in large-scale networks. First, the design approaches currently employed in designing PSTN networks and elements must be carried forward to next-generation networks. These approaches, which allow for real-time recovery and upgrade of networks, are surprisingly complex and require substantial engineering art. Moreover, to support PSTN-level QoS, new algorithms are required to manage the traffic load from packet networks, which exhibit some long-range dependence characteristics of fractal processes.[18] As such, intelligent switches will need to manage emerging protocols such as differentiated services (DiffServ), MPLS, and

RSVP to honor SLAs. Some services, particularly voice, will need to have "guaranteed" fixed bandwidth while others, such as electronic funds transfer, will tolerate variable bandwidth. The network, through policy-based network and service management systems, will need to provide these service-level agreements on demand and be able to adapt to changes in network state and user applications.

The methods for achieving a level of QoS in packet networks that is comparable to that achieved in voice networks will also include sophisticated traffic management algorithms that exploit explicit routing capabilities in the standards-based MPLS networks. These will work in both IP and ATM environments. Moreover, sophisticated, intelligent core and edge switches that have built-in QoS that can classify, queue, and schedule user packets at wire speed and also isolate traffic flows belonging to different users will be central to achieving the QoS vision.

Two papers in this issue address critical issues in QoS in next-generation networks. "QoS and Differentiated Services in a Multiservice Network Environment" by Balakrishnan and Venkateswaran[19] explores technology options for an integrated services infrastructure and analyzes them for their ability to satisfy multiservice requirements. "QoS Routing and Performance in Packet Networks: A Visual Simulation Platform and Case Study" by Funka-Lea et al.[20] describes development of, and analyses carried out using, a general purpose platform to understand and evaluate the performance of new routing and signaling protocols for ATM and IP networks.

Transition

The transition path each carrier will follow to its next generation will be unique, but, as noted above, the transition has already started with overlays of circuit and packet switching and distributed signaling and control. It is useful to observe that in today's voice networks calls are dynamically routed around network bottlenecks using the intelligence of Signaling System 7 (SS7). While today's IP networks use simple, shortest-path route selection and often encounter bottlenecks, it is clear that the evolved packet networks will make use of increasingly intelligent algorithms to meet the application requirements. The specific transition will be

dependent on a carrier's present network, the specific regulatory environment, and the carrier's targeted service offerings. New or evolved network elements will be introduced to support new and existing voice, data, and multimedia services. Simultaneously, the resulting network must be positioned for predictable technology advances, uncertain protocol evolution, and uncertain advanced network services innovation.

Several papers in this issue deal with the critical issues of transition. "Voice and Multiservice Network Design over ATM and IP Networks" by Chen, Das, and Tang[21] describes design methodology, network architecture considerations, and design results from several specific integrated IP/ATM and IP network design studies. "Designing IP Networks with Performance Guarantees" by Benmohamed et al.[22] reports on algorithms for network design and capacity optimization for IP networks incorporating next-generation routers such as the Lucent Technologies' PacketStar IP switch in an embedded network of legacy routers. Finally, "Prototcols, Performance, and Controls for Voice over Wide Area Packet Networks" by Doshi et al.[23] provides a comparative evaluation of alternative protocols, architectures, and controls for long distance voice service over ATM or IP based packet backbones, which may also support many other services over the same infrastructure.

Summary

Dramatic changes in network services, technology, and regulation are creating a golden era of network innovation that is leading to converged, next-generation networks of networks. Much is certain in that evolution. The networks will have:

- A shared, packet-based, optical core network using DWDM optical transport with optical add/drop and multiplexing.
- A variety of broadband access mechanisms—wired and wireless—for all types of customers. Higher-speed access will eliminate the access bottleneck for advanced services.
- Access to traditional and new advanced network services provided by interworking to the PSTN and to new open applications platforms.
- A new style of network management based on

active directories and policy managers.

Simultaneously, many challenges and uncertainties remain. It is crucial that next-generation networks be positioned for expected, yet uncertain advances in these areas:

- Provision of high levels of QoS to support packet switching and network management,
- Evolution and market acceptance of IP and ATM, and
- Rapid service creation.

Finally, the transitioning to high QoS, converged networks is key. Transition must simultaneously position for predictable technology advances, uncertain protocol evolution, and uncertain advanced network services innovation. With these technology advances, it is Lucent's vision that the next-generation networks will be service centric and also that they will bring significant value to both the service providers and the customers by:

- Bringing the reliability and security of voice networks to data networks,
- Reducing the complexity of network management by building intelligence into the products and providing a higher level of network management,
- Supporting a broad range of network servers that enable high-value applications and network services, and
- Offering the best path to a converged voice/data packet network.

This issue of the *Bell Labs Technical Journal* provides a view of the Lucent's Network Vision with a focus on packet networking including network architecture and planning, intelligent packet switching, packet voice, QoS, VPNs, and network and service management. Subsequent issues will provide in-depth treatment of optical networking, services and access, and wireless networking as aspects of Lucent's vision for next-generation networking.

References

1. C. M. Christensen, The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail, Harvard Business School Press, Boston, 1997.
2. D. Mitra, J. A. Morrison, and K. G. Ramakrishnan, "VPN DESIGNER: A Tool for Design of Multi-service Virtual Private Networks," *Bell Labs*

Tech. J., Vol. 3, No. 4, Oct.–Dec. 1998, pp. 15–31.

3.   K. W. Fendick, V. P. Kumar, T. V. Lakshman, and D. Stiliadis, "The PacketStar™ 6400 IP Switch—An IP Switch for the Converged Network," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 32–47.

4.   W. Weiss, "QoS with Differentiated Services," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 48–62.

5.   B. T. Doshi, S. Dravida, P. Harshavardhana, and M. A. Qureshi, "A Comparison of Next-Generation IP-Centric Transport Architectures," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 63–85.

6.   J. M. Fossaceca, J. D. Sandoz, and P. Winterbottom, "The PathStar™ Access Server: Facilitating Carrier-Scale Packet Telephony," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 86–102.

7.   K. K. Whang, W. Fink, B. Krishnamurthy, I-F. Lin, D. M. Rouse, and H. V. Sorensen, "Voice over PacketStar™ Gateway Solution for Service Provider Networks," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 103–123.

8.   J. Kozik, W. A. Montgomery, and J. J. Stanaway, Jr., "Voice Services in Next-Generation Networks: The Evolution of the Intelligent Network and Its Role in Generating New Revenue Opportunities," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 124–143.

9.   D. O. Keck, P. J. Kuehn, "The Feature and Service Interaction Problem in Telecommunications Systems: A Survey," *IEEE Trans. on Software Engineering,* Vol. 24, No. 10, October 1998, pp. 779–796.

10.   Y. Inoue, D. Guha, H. Berndt, " The TINA Consortium," *IEEE Commun.,* Vol. 36, No. 9, September 1998, pp. 130–136.

11.   H. G. Schulzrinne and J. D. Rosenberg, "The Session Initiation Protocol: Providing Advanced Telephony Services Across the Internet," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 144–160.

12.   P. B. Busschbach, "Toward QoS-Capable Virtual Private Networks," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 161–175.

13.   M. S. Chambers, H. Kaur, T. G. Lyons, and B. P. Murphy, "Voice over ATM," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 176–191.

14.   P. G. A. Sijben and L. F. A. Spergel, "Bridging the Gap to IP Telephony," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 192–107.

15.   "Overview of TMN Recommendations," *ITU-T Recommendation M.3000 Series,* ITU-T Study Group 4, Oct. 1994.

16.   "TMN Management Services and Telecommunications Managed Areas: Overview," *Recommendation M.3200, ITU-T Recommendation M.3000 Series,* ITU-T Study Group 4, Apr. 1997.

17.   G. Lin. J. D. Price, and T. K. Srinivas, "Network Information Models and OneVision® Architecture," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 208–221.

18.   W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic," *IEEE/ACM Trans. On Networking,* Vol. 2, No. 1, Feb. 1994, pp. 1–15.

19.   M. Balakrishnan and R. Venkateswaran, "QoS and Differentiated Services in a Multiservice Network Environment," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 222–239.

20.   C. A. Funka-Lea, C. L. Janczewski, W. C. Lau, R. Nagarajan, Y-T. Wang, and Z-L. Xin, "QoS Routing and Performance in Packet Networks: A Visual Simulation Platform and Case Study," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 240–254.

21.   Z. Chen, K. P. Das, and B. Y. C. Tang, "Voice and Multiservice Network Design over ATM and IP Networks," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 255–272.

22.   L. M. Benmohamed, S. Dravida, P. Harshavardhana, W. C. Lau, and A. K. Mittal, "Designing IP Networks with Performance Guarantees," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 273–296.

23.   B. T. Doshi, E. J Hernandez-Valencia, K. Sriram, Y-T. Wang, and O-C. Yue, "Protocols, Performance, and Controls for Voice over Wide Area Packet Networks," *Bell Labs Tech. J.,* Vol. 3, No. 4, Oct.–Dec. 1998, pp. 297–337.

DOUGLAS C. DOWDEN is Technology, Architecture and Platforms Vice President for the Switching and Access Solutions (SAS) Group of Lucent Technologies. He is responsible for driving product architectures and platforms in support of the SAS product teams using advanced technologies and standards. He also serves as the SAS Chief Technology Officer. Based in Naperville, Illinois, he supports an organization located in New Jersey and The Netherlands, as well as in Illinois. Mr. Dowden, a Bell Labs Fellow, holds a B.S. in computer engineering from the University of Illinois at Urbana-Champaign and an M.S. in electrical engineering and computer science from the University of California at Berkeley.

RICHARD D. GITLIN is Chief Technical Officer and Data Networking Vice President of the Data Networking Systems Group of Lucent Technologies. He is responsible for identification and development of core data networking technologies, leading product line architecture work, and formulating and promulgating standards and intellectual property positions. Prior to his current assignment, he was Research Vice President of the Communications Sciences Research Division of Bell Labs, where he led research in high-speed networking, speech, image and message processing, and wireless technology and systems. He has conducted and led pioneering research and development in several areas of digital communications and networking. Well-known for his personal pioneering work in V.32/V.34 modems, for the co-invention of DSL technology, for leading the effort to provide advanced ATM and IP systems and devices, and for his leadership in establishing third-generation broadband wireless systems, he holds a doctorate in engineering science from Columbia University in New York City and is both a Bell Labs and an IEEE Fellow. Dr. Gitlin, who holds 43 patents, has published over 100 papers, winning three best paper awards, and has co-authored a text on data communications as well. He is a member of the Computer Science and Engineering Advisory Committee of the National Science Foundation.

ROBERT L. MARTIN, Chief Technology Officer at Bell Labs, has oversight responsibility for all Lucent Technologies technology strategy and core competency activities. His career, which includes tenure in several other telecommunications organizations/companies as well as at Bell Labs, encompasses a rich and varied set of appointments in software systems. He holds a B.S. degree in electrical engineering from Brown University in Providence, Rhode Island, and M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in Cambridge.◆