CONFERENCE ANALYSIS

IN-DEPTH REPORTS ON LEADING IT CONFERENCES

NETWORLD + INTEROP
Part 1: Network Convergence

Presented by ZD Comdex and Forums
May 11-14, 1999
Las Vegas, Nevada

TABLE OF CONTENTS

Introduction

Networld+Interop (N+I) began as two separate conferences. Interop showcased diverse systems that could interoperate by adhering to the open TCP/IP standards, and Networld was devoted to Novell's proprietary local area networks. The two have merged to form a large, important networking conference. While many conferences focus on showing and selling products, N+I has continued the Interop tradition of bringing IT professionals a significant, high-quality technical program with tutorials, sessions on emerging technology and standards, industry trends, etc. If you attend N+I, plan on attending sessions, not just visiting booths and cocktail parties.

Because of the size of N+I and the breadth of the program, we are devoting two CAN issues to it. The keynotes, sessions and products on the show floor reinforced two themes this year: (1) the convergence of voice, video and data networks (2) and the recognition that supporting and building lasting relationships with customers, not selling to them, may be the killer application for electronic commerce. In this issue, we highlight the remarks of four keynote speakers who centered on convergence. William T. Esrey, of Sprint represented the telephone industry. The others, John Chambers and Judy Estrin of Cisco Systems, and Jeanette Symons of Ascend Communications, represented the IP data community. The next issue will address customer relationships.

The convergence of data and real-time audio and video traffic will require protocols for guaranteed end-to-end quality of service (QOS) across diverse networks. A practical session by David Passmore of NetReference, put QOS in perspective. Convergence will also require ubiquitous, fast last-mile connectivity, which competition will encourage. We covered a session on three emerging last-mile alternatives: high-speed microwave, passive optical networks to the home or business, and low-earth orbiting satellites. We also ventured onto the show floor, where we saw emerging products and standards for high-speed remote access and storage area networks (SANs). We were also briefed on the convergence plans and products of Lucent, a leading telephone equipment manufacturer.

Keynote 1: Network Convergence and ION

William T. Esrey, Chairman and CEO, Sprint

Convergence of voice, data and video over IP networks was a major theme at N+I this year, and William Esrey's presentation focused on using Sprint's Integrated On-Demand IP Network (ION). He also emphasized the advantages of outsourcing networks.

Esrey began with anecdotes about the consequences — some minor, others major — for people not current on computer and networking technology. He told of how one user repeatedly inserted her credit card into a computer's floppy drive when a Web site she was visiting asked her for a credit card. On a more serious note, he referred to a *USA Today* news story about CEOs being forced out of their jobs because their companies had failed to develop an effective Internet strategy.

He stated that we are transitioning to integrated networks, which has been made necessary by a knowledge explosion and the coming "Tsunami" of electronic commerce. Esrey feels electronic commerce is more than just another window into the marketplace. Rather, it has the potential to become a marketplace unto itself. Merchandising, transaction processing, delivery and customer service all need new solutions. IT will go straight to the heart of defining, penetrating and serving markets.

This implies a new role for CIOs and network managers because the market will exist in their converged networks. Network managers and CIOs must add value to information, turning data into intelligence, the currency of the new networked economy. The CIO will have to invent new applications and do strategic redesign of company processes and procedures. Those who add the most value will reap the most rewards. The traditional duties of the network manager are going the way of the keypunch operator. In the future, the IT team will not be seen as a cost center, but as a business center. Companies must transcend the technology to focus on adding value. They must devote less time and resources to running networks and concentrate on using the network to run the business. In other words, companies should outsource their networks to companies such as Sprint.

According to Esrey, outsourcing is driven by three conditions. First, companies cannot keep up with the technologies in their core business while staying current on the latest support technologies. Outsourcing functional expertise is a competitive necessity because it frees time and resources for core technology. The need for highly qualified people is the second motivator. The third is technological obsolescence. Buying hardware and software involves a large risk because technology becomes obsolete quickly. Companies must look beyond the wiring to the wisdom.

Sprint no longer sees itself competing with traditional telecom companies or as a product or service company. Rather, it sees itself as an integrated communication company, concentrating on developing solutions to problems of collaboration, not communications. In March 1997, Sprint's engineers recommended that the company collapse voice and data onto a single integrated on-demand network, now called ION. According to Esrey, ION will deliver unlimited bandwidth on demand and will always be state of the art. It uncorks the flow of knowledge, he said.

ION will integrate all customer premises traffic into an integrated services hub for local and long distance calls and local area and wide area network traffic. This is fed over a high-speed link to the network, enabling many people to work simultaneously. Sprint will be responsible for the sales, installation, maintenance and upgrade of the hub and the network. Companies will be able to expand or contract their bandwidth commitments on demand at any time from their desktops, with their billing rate changing accordingly. For most business, Esrey sees this as yielding savings from 5 percent to 15 percent in the short term and 25 percent in the long term. The network will remain state of the art, according to Esrey. After introducing ION to us, Esrey and a colleague in Kansas City conducted a demonstration.

The Kansas City center had a network node similar to those Sprint has already installed in 26 cities around the country. It consisted of a large asynchronous transfer mode (ATM) switch for service to large businesses and a digital subscriber line (DSL) access multiplexor to provide connectivity to smaller businesses and homes. For the demonstration, Esrey and his colleague had all the makings of a home, small office and large office. Each of these had standard PCs, telephones, etc. connected to a prototype of the integrated services hub.

They began with the home office, demonstrating simultaneous telephony (two toll-quality calls), small-window videoconferencing with a shared whiteboard, Internet surfing, and downloading from the Internet — all using ATM over an asymmetric digital subscriber line (ADSL) (6mb/s to 8mb/s downstream and 1mb/s upstream) connected to the access multiplexor in the simulated central office. The integrated services hub handled all protocol conversion between ATM and the devices inside the office. The single ATM connection allowed bandwidth sharing, dynamic bandwidth allocation and different classes of service for different protocols. The user could control bandwidth and service classes using a single GUI program. The same management program controlled services such as caller ID and call forwarding and could display current billing information. In the large business demonstration they showed virtual connectivity to the corporate LAN. They also spoke of remote security cameras in the home, but did not demonstrate them.

Esrey's demonstration was smooth and impressive. An ION connection in the home would be most welcome, but there was no mention of cost or rollout schedules, and subsequent queries did not yield answers in time for this publication. We were also unable to obtain details on the customer premises equipment. Users will demand competitive pricing if they are to adopt a proprietary solution.

Keynote 2: Convergence for Business, Education and Entertainment

John Chambers,(Live Via Satellite), President and Chief Executive Officer, Cisco Systems, Inc.

Judith Estrin, Senior Vice President and Chief Technology Officer, Cisco Systems, Inc.

This was a double-header keynote. Chambers opened the session with a remote video presentation, promising a new world in which IP-based internetworking would revolutionize business, education and entertainment. According to Chambers, it will change every aspect of our lives — the way we work, play and learn.

He outlined several stages of Internet development. In its early stages, the Internet was used by the academic and research community. Next, businesses began to understand the importance of networking (we are still in the early stages of that understanding), which was when the third, consumer-driven wave began. Chambers compared the Internet revolution to the industrial revolution, predicting globalization and a leveling of the playing field between large and small companies. Like Esrey, he sees the IS role evolving in three phases: from a necessary expense item to a productivity tool and finally to a strategic resource tied to the very survival of the company. Therefore, CEOs are spending increasing time and money on IS.

Government leaders also recognize the IP revolution. Chambers has talked with nearly every government leader in Asia during the last year, and all of them agree on the strategic importance of the Internet to their economy and society.

Telephone companies also realize they are facing startling change. Telephony, the service from which telephone companies derive the majority of their revenue, is becoming a commodity that they may one day provide for free. The carriers realize there will eventually be only one IP network that delivers all traffic. This change will be driven by cost and a need for open standards. Voice and video will ride on the data network, and the convergence is already taking place at an "Internet pace."

Cisco has seen the value of networking in its own operations. Cisco is a virtual company, handling over 70 percent of its service and taking over 75 percent of its orders over the Internet. It has 24 plants around the world, only two of which are Cisco's. It has experienced dramatic productivity increases and is the fastest growing most profitable company in history, with a market capitalization of around $180 billion. The Internet will change business operations forever. Being customer driven will become mandatory, and companies will have to have their finger on the customer's pulse to succeed. Similarly, employee empowerment requires access to information. The company culture must change.

According to Chambers, the message is simple -— are you ready? Is the education system in your state or country ready for the networked age? Is your infrastructure sufficient? Do government leaders understand the importance of networking? Do businesses? Are you ready for the second industrial revolution?

Estrin began with an historical review of the evolution of networking, beginning with the connection of terminals to mainframes. Next came the connection of PCs to LANs for e-mail and file and print servers, followed by the killer app — the Web. Initially, the Web allowed us to search and display information, and conduct customer service, but it now allows us to completely reengineer business practices. We are moving toward intranet and Internet ubiquity with new workflows, methods of collaboration, communities of common interest and electronic marketplaces. The Web is no longer about connecting computers, but connecting people.

This has all happened very quickly, said Estrin. Reaching 50 million users took radio 38 years, PCs 16 years, television 13, and the Web 4 years. (But when do you begin counting for the Internet?) We are still in the early stages; to achieve the full impact of the Internet, we need ubiquitous connectivity. Full ubiquity will require greater bandwidth and reach, data, voice and video convergence, as well as increased network intelligence. Estrin discussed all three.

In considering access, we must view mobile and fixed users as equally important. The Net phone will probably be the most common access device. At the network core, wave-division multiplexing and other optical technology is yielding phenomenal improvement in cost and performance, leading to a flattening of the backbone.

According to Estrin, we need to bring the network to all consumers' homes. Today the options are DSL, CATV and wireless. (Some future contenders are discussed below.) She is often asked which will win, and her answer is "either all or none" because competition among all of them will drive technology and acceptance. The winner in a particular situation will depend on regulatory structures and existing infrastructure more than technology. The same services must run over all media so people can move freely.

We must also drive high-speed access to all business desktops. Enterprise network speed has risen by a factor of 10 every two years since 1995 as we moved from 10 baseT to gigabit Ethernet. We will stay on track with 10 gigabit Ethernet in 2001. At the same time, network intelligence has increased as we have progressed from shared media to multilayer switches and routers. To provide this intelligence at increasing wire speeds, we will see a trend toward substituting custom integrated circuits for software in switches and routers.

Estrin shifted her emphasis from bandwidth and reach to convergence, pointing out that today's circuit-switched public infrastructure was designed for voice traffic and voice communication patterns (short call duration). Data traffic now exceeds voice traffic and is growing at a much faster rate. That means we need an IP/ATM infrastructure, which is already being rolled out. The second phase (in three to five years) will converge content, especially video, bringing entertainment onto the Internet. The telecommunications industry is moving along the same path as IT departments did 20 years ago, moving from centralized, proprietary networks to distributed, standardized (IP) networks.

We will see the data world (its architecture) and the telecommunication world (its availability and simplicity of operation) merge. Three factors will drive convergence: cost savings (per solution, not only per bit); new functionality from applications which can integrate voice, video and data; and competition. This is not about doing telephony better, but about new applications. Cisco has saved money using the Internet by doing things in new ways, not by cutting costs.

Reliability and availability are critical, said Estrin. It is a myth that there is something inherently more reliable about a circuit-switched telephone network than an IP network. Decentralized IP networks route around problems. They also scale more rapidly and reliably than circuit-switched networks and save on connection setup time, which must be amortized over the duration of a session. The unreliability of today's Internet is transitional because the technology and management techniques are new and changing rapidly. On the contrary, the telephone company has had many years to learn.

Some challenges we face during this transition are lowering router convergence times so the network quickly reconfigures when a node fails. We must also design services optimized for packet networks. Rather than building circuits on top of packet networks, we need to increase packet QOS and security. We must also maintain connectivity with legacy systems, and ATM plays a pivotal role here.

The evolution to this new world is under way. Today, IP is being integrated into the system as a bypass or data overload facility. The next step is integrating voice over IP (VOIP) more heavily. There is a similar transition going on in the enterprise. Over the next 12 to 18 months we will see LAN servers assuming the role of the PBX (Private Branch eXchange). Call center customer support will be the killer application driving this transition. (Cisco has acquired call center companies.) Web-based self-help will be integrated with human-based telephone help. The helpers can be specialists at home or at any location.

Increasing network intelligence both at the application and lower levels was Estrin's third major topic. Information appliances will push the Internet to new users who need ease of use and service beyond what we have known in the past. PCs will continue as critical network clients, but IP phones, electronic books, units in cars, pagers, Web phones, etc. will be even more common. These will be “thin” clients — less intelligent than PCs.

Heterogeneous thin clients will demand more intelligence of the network. Hosted applications for companies (or households) that choose to outsource will also demand more network intelligence.

The network must become both application aware and user aware, connecting users to services, not terminals to computers. We will need more complexity in the network to provide easy-to-use and easy-to-manage applications. End-to-end services for QOS, security, caching, multicast, voice and video-enabling, SNA evolution, and operations will be required. Estrin's analogy of Federal Express pointed out that the company went beyond the capabilities of the post office or UPS by adding a new level of value, which made transportation and logistics an integral part of the way people managed inventory. FedEx added services like time-certain delivery, classes of service, and package tracking. Networks will follow the same path, offering new services that can be automatically managed through parameterized policies. We also need comprehensive registration and directory servers. We cannot continue to rely on armies of network managers.

Estrin went over examples of new services we need:

➙ SNA services for interoperation with legacy systems

➙ Policy management for networks and applications

➙ Multicast to facilitate data mirroring, database replication, distributed operating systems, software and information update, push applications, and audio and video conferencing and broadcast

➙ Video-enabling services for policy control and content management

➙ Automatic and transparent caching services that do necessary translation, reformatting and compression

➙ Security services for authentication, perimeter control, intrusion detection, encryption, etc.

➙ QOS from end to end

➙ Voice-enabling services in which call signaling and services control (like 800 numbers) are in layers above IP

She wrapped up the talk by reiterating that policy-based management and comprehensive directories are necessary if we are to be able to manage these new services. She also pointed out that while the technology is exciting, issues such as corporate culture, government regulation, and privacy policy are also critical. Estrin spoke both generally and specifically, describing the direction Cisco's products are taking. The company is betting heavily on its vision being accurate.

Keynote 3: The Path to Core Network Convergence

Jeanette Symons, Chief Technology Officer, Co-founder, Ascend Communications Inc.

Symons began her presentation off-stage, with a humorous video creating the illusion that she was delayed by a late flight and traffic. She "used" a cell phone, laptop, video camera, PowerPoint slides, etc. to "start" her presentation without the benefit of a converged network.

When she finally "arrived" at the convention center, she spoke on network convergence — where we are, how we got here, and where we are going. The first generation Internet was a collection of routers, she said, but we found that Internet service providers (ISPs) could not manage the networks as they grew larger. The next step was to install layer-two switches in the large ISPs' networks. Initially these were Frame Relay switches, now they are ATM switches, and will become optical switches. ISPs have done this in order to be able to manage and control their networks and provide service-level agreements (SLAs) guaranteeing QOS to their customers. This sufficed for e-mail and Web surfing, but will not work for converged networks.

Different groups want convergence for different reasons. The incumbent local exchange carriers want convergence because they believe data traffic far outstrips voice, and they believe the prediction that voice traffic will be so minimal as to be free. They want to incorporate data into their networks. Data networks say they already know how to carry large amounts of data, and they want to pick up the revenue from the relatively small voice volume. Users just want one network so they can integrate their applications and lower their costs. They do not care who does it or how they manage the transition.

The market has decided that the core of the backbone will be IP, not ATM or time-division multiplexing, but we have to figure out what it will take to get there, said Symons. This converged IP network must have the characteristics of today's switched telephone network: 100% availability, an addressing scheme that reaches everyone in the world, and nearly instantaneous setup with no incorrect connections. The average Internet user stays on line for 45 minutes (up from 15 minutes two years ago), but the average voice call is only two to three minutes, making call accounting a problem. The telephone billing system is also secure. The phone network also dedicates a 64kb/s channel with low latency, jitter and error rates to each conversation. It sounds clear, there are no detectable delays, speed up/slow downs or skipped sounds. This service is provided to hundreds of millions of callers every minute.

Today's IP routers make their best effort at packet delivery, but make no service guarantees. So we (the internetworking community) have invented varying types of service (TOS), and put TOS bits in packet headers. This enables us to prioritize data. There are still no guarantees — we continue to do best-effort routing, but some applications, such as voice, get higher priority than others, such as e-mail. Within a LAN or on an enterprise we can provide sufficient bandwidth that TOS works for our applications. However, in a multivendor nationwide or global network, we do not have that degree of control. We need standards that are accepted by all equipment vendors, carriers and governing bodies, according to Symons.

Would you use today's Internet as your sole voice solution? No. It could not deal with hundreds of millions of concurrent calls and still handle all setups in under, say, 1/10 second. The network must be bandwidth-efficient so billing and performance information can be gathered in the background while you make a call, download some data and while your kids are watching a movie.

IP could offer all this today by overengineering, by providing massive bandwidth and redundant nodes and connections, but that would be very expensive — requiring perhaps six times the optimal bandwidth. Another alternative would be microcontrol over bandwidth allocation, call setup, call teardown, etc. within the network. However, neither of these options are good for the long run — they are costly, single vendor options. Our routers and switches will have to change, said Symons. We need new IP protocols.

Without specifying what these protocols are, let's assume they have been invented and agreed upon. How long would convergence take? Let's say Ascend or some other vendor had it engineered into its product line within one year, and it found a carrier willing to deploy it. It would take another year of operating experience to get the equipment and systems around it fully debugged. How long would it take to get the standard approved? Say another year or two. For illustration, we can assume that other vendors begin designing compatible equipment while the standard is being finalized. Once the competing equipment was ready, another carrier could buy it and rollout an interoperable service. The absolute best case for the start of such a multivendor, interoperable network would be five to six years out. Global coverage would take far longer. In the meantime, we will continue doing data networking and voice over IP in controlled situations and vertical markets.

Once true global convergence is achieved, what will we have? Will we get rid of our telephones and do everything from our PCs? What will the converged network look like? Symons feels we will see new metaphors and applications. That is the purpose of core convergence, not reduced cost of calls.

Who will do this? The router companies might, but they must achieve the scalability, reliability and other characteristics of the voice network. Will it be the traditional telecommunication equipment vendors? They would have to learn to touch the data — to build devices that make decisions based on data content. Who will win? Symons does not know. Will it be one of the big players you see today? Probably — because the carriers need critical mass. The successful little startups will end up inside the large companies. Five or six years from now, we will have a converged network, and that network will have novel metaphors and applications. (She and Estrin both mentioned five years, but after hearing Symons development and rollout scenario, it is hard to imagine that happening on a wide scale so soon.)

Session 1: Quality of Service Strategies for the Enterprise

David Passmore, NetReference, Inc.

Passmore's discussion of QOS strategies for the enterprise covered QOS requirements, alternative QOS technologies, future developments and evaluation criteria, as well as recommendations for the enterprise.

IP networks treat all packets the same, making their best effort to deliver them on a first-come, first-served basis. Passmore addressed the question, "How can the network deliver different service levels to different applications in support of the enterprise's QOS policies?" In other words, how can we build "unfair" networks, giving high-priority traffic preference over low-priority traffic.

Before looking at technology to increase QOS, organizational policies can be used, he noted. For example, a company can ban playing Doom on the network, and can prohibit users from installing applications without approval. Overprovision of bandwidth or reliance on an ISP's contractually guaranteed SLAs can also guarantee QOS.

Passmore categorizes enterprise traffic as follows:

- Routine or bulk transfer, such as Web surfing, e-mail, file transfer and buffered streaming media
- Transaction processing, such as mainframe SNA sessions or protocols and applications subject to timeouts
- Real-time multimedia, such as video and voice
- Network control flows, such as signals that a part of the network is overloaded

Passmore noted that latency (packet delay) is as important as bandwidth for applications like telephony and videoconferencing. Transactions may time out in one or two seconds, and real-time conversations seem stilted if latency exceeds 100 to 150ms. Real-time conversations are also disrupted by jitter (variation in latency). It is increasingly the case that bandwidth is easy, but meeting latency constraints is difficult.

Real-time requirements are most stringent. Speed-of-light delays are roughly 10 ms./1,000 miles, making geo-synchronous satellites unsatisfactory. (Low-earth orbit satellites will not have this problem.) Time for the application to run, as well as compression, routing and decompression must be added to transmission time. A time-consuming buffer is also needed to smooth out jitter. Error rates must also be low. Since there is no time for retransmission, bad packets are simply discarded. (Good systems inject noise in place of discarded packets, making them less noticeable.) Since we do not need retransmission, User Datagram Protocol (UDP) is used in place of TCP, which checks for errors.

Passmore turned to alternatives for achieving QOS. While his focus is on IP networks, he first discussed ATM and Frame Relay to put IP in context. He then covered IP and also Ethernet QOS.

ATM has different service classes and also more intelligent routing than IP. It is mature, and was designed from its inception around QOS. Unlike IP, ATM sets up virtual circuits. In doing so it takes factors such as expected delay into account. It also compiles statistics for SLAs and monitors traffic for violations. Traffic can be shaped at entry and ingress rate controlled. ATM defines five service level categories:

- Constant bit rate (CBR)
- Real-time variable bit rate (rt-VBR)
- Non-real-time variable bit rate (nrt-VBR)
- Available bit rate (ABR)
- Unspecified bit rate (UBR)

CBR is circuit emulation with predictable bandwidth and latency. It is used for telephony, providing toll-quality service, but it does not use bandwidth effectively. With rt-VBR you specify a peak cell rate, sustained cell rate, maximum burst size and cell transfer delay (beyond which packets are dropped) at the time the virtual circuit is set up. It is used for compressed audio or video where there are variable bit rates due, for example, to variable amounts of motion. nrt-VBR is similar to rt-VBR, but cells are not dropped when a cell transfer delay is exceeded. This is useful for transaction processing. ABR makes no latency guarantee, but does guarantee bandwidth. The network can also throttle or provide flow control at the end points. This might be used for IP traffic or bridging LANs together. UBR makes no guarantees and has no feedback mechanism between the endpoints, but it is cheap. It would be fine for e-mail or FTP.

In addition to classes of service, ATM routing protocols help with QOS. Routing is not hop by hop; an entire path is built and analyzed at setup time and congestion is considered. If QOS requirements for bandwidth, latency, jitter and error rate cannot be satisfied, the connection request is refused. IP routers look only at hop count, and may send packets along short congested paths when longer open paths exist. When the network is becoming congested, ATM can slow ingress traffic. Protocols also exist to map ATM in the core to Ethernet LANs, maintaining guaranteed QOS.

Passmore spoke briefly on Frame Relay. Like ATM, Frame Relay is connection oriented. Users pay for a committed information rate (CIR) and the network is guaranteed to deliver at least that much bandwidth. You may be able to burst above the CIR. He cited the example of his own company, which has a CIR of 256kb/s, but often bursts well above that. This results in near-T1 performance at a much lower cost. Applications can also set a discard-eligible bit, indicating which packets should be dropped in the case of congestion, but not everyone uses it. Frame Relay does not perform error checking or retransmission, but TCP does. As with ATM, end points are notified about congestion, and may slow down. Some Frame Relay networks also have proprietary priority systems, but there are no standards at the present time.

With the ATM and Frame Relay contexts set, Passmore turned to IP QOS. Three techniques are used: packet scheduling and prioritization, congestion control, and traffic shaping. In a router or switch, a packet is classified when it arrives (voice, video, transaction, etc.). This may be done by examining TOS bits, source and destination IP addresses or port numbers, destination URL, etc. Once the packet is classified, it is scheduled for transmission using some queue management algorithm. The algorithm may also consider which flow a packet is part of, for example, if it is part of a specific voice conversation. He also pointed out that there is a tradeoff between large buffers (cost of memory and higher jitter) and small buffers (packet loss).

There are interactions between router action and the TCP protocol at the end point. If a router drops packets, TCP will stop transmission temporarily. Since TCP has a slow start algorithm, it starts up slowly at first, wasting bandwidth, then builds back up, perhaps causing another shutdown. This is called the global synchronization problem, and it may yield oscillating waves of traffic.

Passmore reviewed several queuing algorithms, beginning with first in, first out. Next was priority queuing, in which there is a separate queue for each class of packet, and low-priority packets are transmitted only if the high-priority queues are empty. If low-priority packets stall for too long, their applications may time out, sending an error message resulting in retransmission, exacerbating congestion. Fair queuing is a "round robin" scheme in which every packet is treated the same. Weighted, fair queuing uses prioritization, but applies some aging, so old low-priority packets may be transmitted before new high-priority packets. There are many variations on this theme. For example, packets belonging to low-bandwidth flows may get high priority.

Passmore described a congestion-control technique called random early detection (RED) in which the router purposely drops packets in an attempt to fix the global synchronous problem mentioned above. Packets are dropped from each stream at random in order to prevent buffer overflow in the router. There is also a weighted variation of RED, in which fewer high-priority packets are dropped. RED works with TCP, which does error checking, but would have no effect on UDP packets (which are used in many streaming and real-time applications).

He also discussed traffic shaping, in which TCP window sizes (transmission rates) are manipulated to smooth traffic. They emulate a "leaky bucket," keeping output rate constant regardless of the input rate.

Passmore turned from transmission scheduling to packet classification. How does the router know what kind of traffic it is dealing with? An application or user may signal this explicitly or the router can do this implicitly by examining the packet and automatically classifying it.

The resource reservation protocol (RSVP) and Differential Services (DiffServ) are two explicit mechanisms. With RSVP, a server requests that a path with certain guaranteed transmission characteristics be set up across the network, and the network acknowledges when the routers along the path agree to the contract. Unfortunately, RSVP has not scaled well beyond the LAN and constrained enterprise environments. Since few applications use RSVP, some routers can be configured to automatically make RSVP calls for traffic on certain ports. RSVP is supported in Windows 2000.

DiffServ uses an 8-bit precedence field, which has always been in IP packets, but generally has been ignored until now. There is a major effort in the Internet Engineering Task Force to define the use of the precedence bits in a consistent manner. As with RSVP, an application or an edge router can set the DiffServ level. DiffServ scales well, but does not guarantee service, while RSVP scales poorly and guarantees service.

Implicit signaling mechanisms use information that is already in the packet. For example, if you have an Ethernet phone at a certain IP address, you know it should have high priority. This is simple conceptually, but implies a lot of work in setting up tables. The management of this function must be automated for it to succeed. You can also look more deeply into the packet at the port number. For example, you can identify Web traffic because it conventionally uses port 80. There are problems with nonstandard port assignments and applications like H.323 voice, which uses dynamic port assignment for data flows. Few routers can track dynamic port assignment, and those considering voice over IP applications should discuss this with their router vendors. Finally, you can look deeper into packets. For example, looking at the URL within a Web packet would allow you to give high priority to certain customers or to customers who were about to place orders in an electronic commerce application. Any of these options would require automated management.

Having finished his discussion of IP, Passmore turned to Ethernet QOS. The IEEE 802.1p standard defines a user-priority field, which is in a new virtual LAN header defined in IEEE 802.1q. Protocols using the user-priority field would be implemented in layer 2 switches at the edge of the network. If your need for Ethernet QOS is not pressing, you might hold off. Passmore predicts that in a few years, we will have moved to level 3 switching at the edge of the LAN, detecting IP addresses, and eliminating the need for IEEE 802.1p and q.

There is also an Ethernet flow control mechanism defined by the IEEE 802.3x standard. This is similar to x-on/x-off to stop and start flows. It could be used by a switch that is receiving packets at 100mb/s and sending them out at 10mb/s.

Passmore made several specific predictions regarding IP QOS. For one, DiffServ levels will be standardized by the Internet Engineering Task Force (IETF) so all vendors can share the same vocabulary for each level. This is necessary so traffic can be routed appropriately across multivendor networks. This will take a while, but it is moving forward since the equipment vendors are highly motivated. He also looks forward to Multi-Protocol Label Switching (MPLS), which provides routing in which path choice takes congestion and other factors into account in the manner of ATM. Interoperability between different technologies such as ATM, Frame Relay and RSVP-based IP networks is also being worked on. This promises end-to-end QOS across different types of links.

These are fine, but the most important development will be policy-based networking. There are many QOS mechanisms, and once a system administrator has decided what to deploy, it is necessary to configure the network. Today that means manually turning router features on and off and entering tables of addresses and other parameters the router algorithms uses. This process is error prone and can lead to unforeseen interactions and side effects.

Since manual management is impractical, policy servers will use knowledge of the network configuration, applications, and users to automatically configure routers and switches. All of this information will be stored in an LDAP-compliant directory such as Novell Directory Service or Microsoft Active Directory. There is an effort called Directory-Enabled Networking (DEN) to standardize the network directory. Microsoft and Cisco have led this effort, and Passmore feels it will succeed. This all looks promising, but for now there is no multivendor interoperability, and you must buy all your equipment from one vendor for the time being.

Having surveyed the technological options, Passmore turned to evaluating QOS alternatives. Criteria include the ability to meet the service-level requirements of your applications and sufficient granularity to classify your packets — how many service classes are needed? Can QOS guarantees be defined, fulfilled and enforced? Is a prospective solution easy to administer? (He feels many people will need consultants to set up their policy-server schemas.) Can charge-backs be monitored? (He feels we will probably not see usage-based billing, but unlimited service at several differentially priced priority levels.)

Passmore made definitive recommendations. For LANs without voice and video traffic, one should increase bandwidth using switched Fast and Gigabit Ethernet, and not bother with QOS. For those planning voice over IP or PBX replacement on LANs, minimal prioritization in all switches and routers is recommended.

On a large WAN (over 20 nodes) with voice or video over IP, you need ATM today, said Passamore. That is what service providers use. For smaller WANs, real-time audio and video can be handled with weighted fair queuing and RED in the routers, and RSVP signaling. If all you have is transaction traffic, turning on weighted fair queuing with implicit signaling should suffice. If you do a lot of bulk data transfer, you should also be doing some rate shaping at the edge to prevent flooding the network. If you have none of those requirements, stick with standard best-effort routing.

In his closing remarks, Passmore repeated that throwing bandwidth at a problem is often better than QOS. He also pointed out that after rejecting ATM for its complexity, we are slowly moving IP in the same direction. Ease of administration is the key to making QOS work, and the vendors have a long way to go. Even when the vendors solve all their problems, you must still decide who gets which level of service, which may be a sticky political question. Finally, QOS will really become a major consideration when voice over IP becomes popular and people begin replacing their PBXs.

Passmore left us with some practical, concrete recommendations. Many of his predictions were corroborated by Estrin in her keynote address. They agree that we can look forward to

DiffServ in the WAN; more custom ICs in routers; directory-driven, policy-based QOS, etc. This is not surprising since Passmore has doubtlessly been briefed on Cisco's product plans. Nevertheless, it is a clear indication of where we are headed. For further information on QOS, visit the QOS Forum at (www.qosforum.com).

Session 2: Last-Mile Technologies

Francois Vigneron, Product Line Manager for Fixed Wireless Access, Alcatel.

Stephen Denny, Senior Product Planner, Fujitsu Network Communications, Inc.

Todd Kaloudis, Manager of Technology Strategy, Teledesic.

Cable TV and xDSL are currently used for high-speed Internet access in homes in prosperous countries, such as the United States and Singapore, and most businesses have high-speed leased lines. This session presented three emerging last-mile alternatives: Local Multipoint Distribution System (LMDS) microwave links, passive optical networks to the home or business, and low-earth orbiting satellites.

Vigneron spoke on LMDS, a point-to-multipoint microwave alternative for the last mile. The initial focus for LMDS will be small- and medium-sized businesses, with residential service at a later time. LMDS offerings will run from 200kb/s to 155mb/s. Last year, the Federal Communication Commission (FCC) auctioned spectrum, and two competing licenses were awarded in each of 493 areas in the United States. The majority of these were to competitive local exchange carriers, and only 1 percent went to existing regional Bell operating companies (RBOCs). The FCC raised $600 million, and these new licensees will soon be in competition with the RBOCs.

Only 10 percent of Alcatel’s targeted office buildings and business parks is currently reached by fiber. A high-speed ATM link over LMDS might connect an operator’s base station to an end user or ISP. Downstream connections might be voice, DSL, Ethernet, T1, fractional T1, etc. Vigneron expects that most subscriber services will be asymmetrical, with more downstream than upstream bandwidth.

Customer premises equipment will consist of a 10- to 12-inch parabolic antenna with a transceiver on the top or side of the building and a cable to an equipment rack in the building. The base station covers a 4-kilometer-radius circular area. The coverage is divided into four 90-degree sectors (polarization alternates between horizontal and vertical in adjacent sectors), each accommodating about 1,000 customer premises. When a base station becomes saturated, capacity may be expanded by adding carrier frequencies. ATM is used within the network — like Sprint ION without the wires. Alcatel will serve two types of customer at first. Large customers will have point-to-point links at speeds up to 155mb/s, and smaller customers will have point-to-multipoint links at speeds from 200k to 8mb/s.

Problems with LMDS include the necessity for line-of-site paths between antennae, which may be a problem in congested downtown areas or rough terrain, and may experience interference due to rain. The former may be mitigated somewhat by cell overlap, and Alcatel will deploy systems which are tuned appropriately for the expected rainfall in the region. The goal is 99.99 percent availability. Vigneron concluded by stating that LMDS was a viable alternative for those wishing to compete with existing telephone companies. He also pointed out that subscribers use bandwidth only on demand, saving money and lowering operator cost by minimizing spectrum consumption.

Won't we eventually all have fiber running to our homes? Are cable modems, ADSL and other solutions interim steps to fiber in our homes? Denny, of Fujitsu Network Communications, Inc., feels escalating bandwidth demands will ultimately justify an all-optical network. Passive optical networks (PON) technology has been available for around 10 years, but ATM PONs are a new offering. Passive optical splitters have no electronics and require no power. They divide the bandwidth on a cable between many users, fanning out to connect many endpoints.

The basic architecture would connect an optical line-termination point in a central office to customer premises equipment in the home. This would be an ATM link with a total capacity of 155mb/s. It could be up to 12 miles long (copper lines are limited to approximately three miles) and split into up to 32 sublinks to individual homes in a neighborhood. This architecture would accommodate voice, data and video. Denny showed a video illustrating a deployment in which a line was first run to a shopping center then split to serve two businesses. Later it was split further to add homes in the area.

Different frequencies are used for the uplink and downlink data, and the splitter is an off-the-shelf device that has been used for quite some time. A standards group called the Full Service Access Network is developing standards for this technology and has submitted a specification to the International Telecommunications Union (ITU). That has become ITU Recommendation G.983.

Fiber to the home has several advantages over copper. Signals are immune to the crosstalk found in copper conductors in the same binder group, there is no influence from lightning, power lines or other sources of electromagnetic radiation. Another key advantage is the large central office serving area enabled by 12-mile long links. (Denny predicted that even that limitation would be increased.) Long cable runs and a passive system mean reduced investment and maintenance cost. The final benefit is the increased bandwidth over copper.

Denny noted that there is already considerable fiber in the ground to neighborhoods, but it is not run to homes. Fiber installation cost is 10 times the cost of material, and rehabilitation trenching is four times the cost of trenching in a new building, therefore fiber should be dropped in all open trenches. If you need to deliver video, have exhausted your copper lines or bandwidth, upgrade existing service, build out to new communities, or want to cut maintenance costs, you should install fiber — it is the preferred solution for the future.

In late 2003, Teledesic will begin offering global satellite service that looks and feels like a managed terrestrial service. The network will offer the following features, which Kaloudis feels will be necessary to be competitive in the future:

- Global presence
- High reliability

- Low latency for synchronous applications
- Guaranteed quality of service
- Open standards
- An easily managed, multifunction network
- Bandwidth on demand

The Teledesic network will be based on a constellation of 288 low-earth orbit satellites. (See www.teledesic.com for details and animated depictions of the network.) The satellites will form a high-speed backbone connecting both large end users and service providers. Since the satellites are in low-earth orbit, and the optical links between them are very fast, latency between any two end points will be acceptable for telephony and videoconferencing.

Kaloudis stated that data rates will begin at 2mb/s uplink and 64mb/s downlink for a low-end system. Bandwidth is on demand and may be turned on and off in between 50ms and 100ms. Latency is projected to be under 75ms in the United States and under 85ms on transatlantic routes. Communication will be hampered by rain, but Teledesic is projecting 99.9 percent availability in the rainiest parts of the United States and a packet loss rate below 1 per million.

Teledesic will achieve this performance by over-designing the network. The orbiting routers will be connected with high-speed links, and spare satellites will provide alternative routes in the case of failure. Even if the spares were down, an outage would be brief due to the rapid motion of the satellites relative to the surface of the Earth.

As protocols evolve, Teledesic will make changes by altering programs at a service adaptation layer between the customer premises equipment and the network interface. The architecture of the orbiting network is flexible with many hooks to accommodate service adaptation. It offers bandwidth reservation, queue management at ingress and egress points, preemption for high-priority traffic, and congestion management to ensure graceful downgrade. All of this is under program control at the adaptation layer. Kaloudis is confident that future QOS protocols like DiffServ will be easy to implement.

Customer premises equipment will have two GPS-equipped, steerable dishes. One will track the current satellite, the other the next one coming. These will connect to an outdoor box connected via cable to an indoor box, which is on a standard LAN. The company will offer three configurations at first:

Teledisic's Network Offerings

	Standard	Enterprise	Aggregator
Uplink speed	2.048 mb/s	N x 2.048 mb/s	N x 51.84 mb/s
Downlink rate	64 mb/s	64 mb/s	N x 51.84 mb/s
Availability	99.9%	99.9%	99.99%
Aperture size	45-60 cm.	45-60 cm.	1.8 m

The customer premises equipment for the standard service will cost in the neighborhood of $4,000 to $5,000. Teledesic is also developing specialized ground-station equipment to be used on moving ships.

Kaloudis feels the network will be used for contingency bandwidth, extending the reach of a current service provider, and transportable network endpoints (for example at a trade show or off-shore drilling site). It will also be useful in developing nations. (For a discussion of Teledesic's potential in developing nations see som.csudh.edu/fac/lpress/articles/villages.htm)

On the Show Floor

Networld+Interop formed when the two conferences merged. From its inception, Interop was geared toward demonstrating the interoperability of IP-based products. As such, a show network was set up, and all vendors had to connect to and operate over it. If you could not interoperate, you could not participate. There were also small "hot spots" on the show floor, which were dedicated to interoperability with specific protocols. There were four hot spots this year (we cover two in this issue and two in the next).

The High-Speed Remote Access hot spot featured seven companies interested in secure access for telecommuters and for people working while traveling. They showed DSL modems, cable modems and VPN solutions from carriers using the Internet Protocol Security standard (IPSec). IPSec is a series of guidelines for secure IP traffic. It is based on a series of RFCs, and provides for encryption, authentication, detection of data tampering or unauthorized resending of data, and key management. IPSec defines a security layer between IP and TCP or UDP, or it may be used around TCP/IP in a tunneling application. To learn more, see the IPSec Developers' Forum (www.ip-sec.com) or look for the forthcoming Prentice Hall book *Ipsec : The New Security Standard for the Internet, Intranets, and Virtual Private Networks* by Paul Mockapetis (Editor), Naganand Doraswamy and Dan Harkins. The Interop Bookstore displayed chapters of this book.

The Storage Area Network (SAN) hot spot demonstrated products of 15 companies with fibre channel links, switches, network adapters, disk arrays, tape arrays and operating systems working together to access, backup and manage storage on SANs. They demonstrated any-to-any access, backup and storage management. The communication took place over Fibre Channel links capable of transmission rates between 266mb/s and 4gb/s at distances up to 10 kilometers. Fibre Channel also guarantees packet delivery. Fibre Channel is a family of ANSI standards, and you can get more information from the Fibre Channel Association (www.fibrechannel.com) and the National Committee for Information Processing Standards (www.ncits.org/tc_home/t11.htm). You might also look at *Fibre Channel for Mass Storage* by Ralph H. Thornburgh. This is a Hewlett-Packard Professional Book published by Prentice Hall. It is short and clear, but several chapters are specific to Hewlett-Packard equipment. You might also look for Addison Wesley's forthcoming *Designing Storage Area Networks: A Practical Reference for Implementing Fibre Channel SANs* by Thomas Clark.

Convergence was the theme of several keynotes, and there were many voice over IP products on the show floor. Some were geared to the LAN, others to the distributed enterprise and others to the carrier community. While full public network convergence is many years in the future, it is clear that pockets of VOIP are with us now, and will rapidly expand. A visit with Lucent Technologies illustrated the view from the perspective of a telephone equipment vendor.

Lucent is clearly committed to VOIP. The company and has purchased Ascend Communication and it agrees with the data-centric keynote speakers that packet-switched convergence will occur in public networks. The company feels VOIP has mostly been used by early adopters for toll bypass, particularly in international calling, but large multinationals and enterprises deploying multimedia applications are now adopting it. When carriers are ready, they will begin to outsource these applications for users by offering service-level agreements on virtual private networks. The final step will be ubiquitous multimedia IP over the Internet. Lucent's product line will address all of these markets.

Lucent demonstrated its current enterprise VOIP solutions at N+I, making calls within a simulated organization and over the public telephone network through its gateway. The sound quality was fine, and call management integrated VOIP and regular telephone calls. We were also briefed on Pathstar, an equipment line introduced two weeks before N+I. Pathstar is not oriented toward the carrier, nor the enterprise. The company will sell it to new competitive local exchange carriers and Internet service providers entering the voice market.

Pathstar is offered as an alternative to a circuit switch, integrating class-5 telephony, DSL access, VOIP and IP access and routing in one system. Pathstar systems consist of Access Servers located in central offices and Business Service Exchanges (BSE), located on the customer premises (for example in a business park or office building) with a connection back to the Access Server. Lucent assumes the BSE would be owned and managed by the provider even though it is on the customer's premises. Pathstar offers over 60 telephony features, Internet integration and one-console management with a capital cost saving of 67 percent and an operating cost savings of 70 percent compared to conventional equipment. If Lucent's estimates are correct, it is hard to imagine a new provider going with circuit switches.

Lucent's microelectronics group was also at N+I, and the company announced its phone-on-a-chip. This is a single-chip combining a codec, speech compression, speakerphone echo cancellation, speaker and microphone amplification, Ethernet, an Ethernet repeater (so one jack can serve your phone and PC), universal serial bus, infrared link, keypad controller and LCD display driver! Lucent expects it to bring the cost of an Ethernet telephone down to approximately $150.

About the Reporter

Larry Press is Professor of Computer Information Systems at California State University, Dominguez Hills, and is a contributing editor to both the *Communications of the Association for Computing Machinery* and *OnTheInternet*, the publication of the Internet Society. Mr. Press studies the applications and implications of computer networks and the global diffusion of the Internet.

Giga Information Group Worldwide Contacts

UNITED STATES
Massachusetts
Giga Information Group
One Kendall Square
Building 1400W
Cambridge, MA 02139
Tel: (617) 577-9595
Fax: (617) 577-1649

Giga Information Group
One Longwater Circle
Norwell, MA 02061
Tel: (781) 982-9500
Fax: (781) 878-6650

California
Giga Information Group
3945 Freedom Circle, Suite 720
Santa Clara, CA 95054
Tel: (408) 987-2765
Fax: (408) 492-9823

Connecticut
Giga Information Group
101 Merritt 7, 5th Floor
Norwalk, CT 06851
Tel: (203) 845-6900
Fax: (203) 840-8923

Illinois
Giga Information Group
1300 West Higgins Rd., Suite 116
Park Ridge, IL 60068
Tel: (847) 823-2393
Fax: (847) 823-2394

New York
Giga Information Group
200 West 57th Street, Suite 1208
New York, NY 10019
Tel: (212) 237-2700
Fax: (212) 977-4564

Texas
Giga Information Group
122 West Carpenter Fwy., Suite 535
Irving TX 75039
Tel: (972) 893-5300
Fax: (972) 893-5301

EUROPE
France
Giga Information Group
98 Route de La Reine
92100 Boulogne
France
Tel: 011 (33) 1 48 25 32 00
Fax: 011 (33) 1 48 25 41 93

Germany
Giga Information Group GmbH
Carl-Zeiss-Ring 19-21
85737 Ismaning bei Meunchen
Germany
Tel: 011 (49) 89 9607 830
Fax: 011 (49) 89 9607 8330

Giga Information Group GmbH
Konigsallee 60F
D-40212 Duesseldorf
Germany
Tel: 011 (49) 211 8903 272
Fax: 011 (49) 211 8903 288

Italy
Giga Information Group
Via Mario Pagano, 12
20145 Milano
Italy
Tel: 011 (39) 2 3310 0365
Fax: 011 (39) 2 3310 1750

United Kingdom
Giga Information Group
Arliss Court
24 Clarendon Road
Watford, Herts, WD1 1GG
UK
Tel: 011 (44) 1923 354444
Fax: 011 (44) 1923 354433

Giga Information Group
Castle Hill House
12 Castle Hill

Windsor, Berkshire SL4 1PD
Tel: 011 (44) 1753 831731
Fax: 011 (44) 1753 83199

Conference Analysis Newsletter
Easy Fax-Back Order Form (781) 982-1724

❑ **Free Preview Issue!**
I would like to preview a complimentary issue of Conference Analysis Newsletter! Please e-mail my sample copy to: ____________________

(Or, just send an email to: emurphy@gigaweb.com with "CAN Sample Issue" in the subject line.)

❑ **1 Year - $497.00 U.S.**
Yes, I would like to subscribe to the Conference Analysis Newsletter. I understand that I will have access to a minimum of 24 issues per year. (Note: Please be sure to include your phone and fax numbers and email address so we can contact you with your new password.)

Name: ____________________ Title: ____________________

Company: ____________________ Address: ____________________

____________________ ____________________

City: ____________________ State: __________ Zip: __________

Phone: ____________________ **Fax:** ____________________

E-mail: ____________________

Please charge my: ❑ Master Card ❑ Visa ❑ Diner's Club ❑ American Express

Account #: ____________________

Expiration Date (Month/Year): ____________________

Signature: ____________________

❑ Check enclosed (Please make checks payable to: Giga Information Group).

Mail to: Giga Information Group, c/o Conference Analysis Newsletter, One Longwater Circle, Norwell MA 02061, Fax: (781)-982-1724. All checks must be drawn on a U.S. bank. Please add applicable state sales tax.

For additional information, please call Eileen Murphy at (781) 792-2610
or e-mail: emurphy@gigaweb.com