Volume #5 *CAN #426*

July 17, 1998

by Barbara C. McNurlin & Laurence I. Press

Conference Analysis

In-depth reports on leading IT conferences

Networld/InterOp 98
Presented ZD COMDEX and Forums
May 4–8, 1998
Las Vegas, Nevada

Introduction

This is the second installment of our continuing coverage of the expansive Networld+Interop 98 show in Las Vegas. —Ed.

Table of Contents

Keynote 2: The Next Generation of Broadband Communications

Rob Glaser, Founder and CEO Real Networks

Mr. Glaser covered the following general topics in his keynote:

- Internet & intranet trends
- Communications based computing (applications)
- Technical challenges & requirements

Internet and intranet trends.

One trend is strong growth. There were an estimated 57 million Internet users in 1997, up from 31.4 million in 1996.

Exponential growth and widespread acceptance of the Internet as a viable shopping medium is expected. Revenues from e-commerce transactions topped out at an estimated $707 million in 1996. Estimates for 1997 rounded out at $2.6 billion in e-commerce transactions. In addition, online ad revenue jumped to $906 million in 1997, up from $268 million in 1996.

Another trend is the growth of widespead acceptance of intranets.

Companies that are:	**July, 1996**	**Sept, 1997**
Considering	10%	12%
Evaluating	7%	16%
Using	14%	50%

Connection speeds are also rising, but office connections are much faster, producing a bi-modal environment:

	Oct., 1996	**Oct.., 1997**
14.4 kbps	21%	7%
28.8-56 kbps	61%	70%
ISDN	4%	5%
>= T1	14%	18%

The great majority of PCs are now multimedia-ready. TCP/IP, HTTP, HTML, and other standards are now in place. IP connectivity is pretty well restricted to PCs today, but set top boxes, wireless connections, cable, etc. well greatly extend that.

Communication based computing (applications).

Communication based computing is also with us and growing. We moved from terminal emulation in the late 1980s to LANs and WANs in the early-mid 1990s to LAN/WAN and Internet today. We have now moved beyond text and can assume wide access outside of our work world.

Mr. Glaser foresees three important communications-based applications emerging:

- Desktop conferencing and telephony
- Multimedia email and messaging
- Streaming media delivery
    - one to many live broadcast
    - one to one, or one to many on-demand

Mr. Glaser predicts that multimedia email and messaging will catch on once people become comfortable with leaving video messages (as they did with answering machines), and we get simple tools for editing and correcting. He feels in three or four years multimedia email will be as common as text based email.

Mr. Glaser's company focuses on streaming media, controlling 85 percent of Web sites that host streaming technology. He feels complex ideas are easier to convey and digest when video is used. Streaming media volume is increasing rapidly:

	April, 1997	**April, 1998**
Streaming web pages (thousands)	80	307
Live hours per week (thousands)	40	150
RealPlayer users (millions)	8	22

This is faster than Internet growth, and comparable to the rate of growth of the Web about three years ago.

There is also growth in live, continuous feeds. There are now over 1,000, 24 hour-a-day live audio and video *stations* on the Net. The number of users has also exploded. Over 20 million users with unique email addresses have Real Player software, and that is growing at a rate of 100,000 downloads per day.

The range of applications on the Internet is also broad:

- Broadcasting
    - news/information
    - entertainment
- Marketing and selling
    - advertising
    - product demonstrations
- Teaching and training
- Personal communication

Mr. Glaser singled out news, citing MSNBC research showing that *people average nearly as much time per week reading online news nearly as many hours per week as newspapers*:

News Medium	**Hours/Week**
Broadcast TV	5.7
Cable TV	5.0
Radio	4.5
Newspapers	3.6
Internet/Online	3.5
Magazines	2.4

He demonstrated short video "headlines" in a streaming context using the latest version of their software.

Mr. Glaser said streaming media is being widely used for executive briefings and other external communication on intranets. He showed clips illustrating executive briefings at Boeing and 3Com, each of which was viewed by over 10,000 people. (The Boeing clip had Spanish subtitles, incidentally.)

Mr. Glaser stated that **Boeing is saving about $130,000 per week using streaming technology.** Sales force training is another strong application for streaming media, and he presented a testimonial from MCI and a clip featuring management consultant and author Tom Peters.

Technical challenges and requirements.

Mr. Glaser asked what the challenges were if we were to make such things universal. He mentioned bandwidth requirements, protocols to support isochronous communication, and standards.

Today we are in a bi-modal period for bandwidth—some have high bandwidth now, others are using modems. Enterprises and on-campus LANs will have broadband capability soon, but the ramping up will be slower in wide-area networks.

The installed base of cable modems grew from 100,000 at the end of 1997 to around 200,000 today. At this rate of growth, it will still take a very long time to reach the 100 million households in the US.

The most optimistic surveys he has seen suggest that by the year 2000 only around 15-20 percent of US households will have high speed access via cable modem, Digital Subscriber Line, or satellite. He cited a Jupiter Communications survey showing the following communication rates to US households today and in the year 2000:

Speed	**1998**	**2000**
sub-28.8K	89%	18%
56K	10%	65%
high-speed	1%	17%

High-speed communication to the home will not be ubiquitous in the near future, so they are introducing new data types that stream well—text and audio with synchronized stills (including client side effects and fades).

Applications must also scale and run on mission critical networks. IP multicasting (see the report in this issue) will be important. They also have a proprietary splitter technology which can be used where multicasting is not available. They have deployed splitters in the US and Japan (with MCI and NTT). Local caching is also important.

Customers also expect broadcast-level reliability and performance, and they are now doing dynamic protocol and bandwidth negotiation (before and during sessions) and interpolating to replace lost packets. He showed an example in which a pause for re-buffering was needed using their old technology and the new technology merely degraded gracefully, maintaining continuity.

At the transport layer, Real-Time Streaming Protocol (RTSP) is now an IETF standard (RFC 2326). At the application layer, we have a newer standard: the Synchronized Multimedia Integration Language (SMIL). Mr. Glaser referred to RTSP as the "HTTP" of streaming media and SMIL as the "HTML" of streaming media. SMIL is now a W3C proposal, and it allows for synchronization of all types of streaming media—video, audio, text, images, etc. They support both RTSP and SMIL in their latest software.

They feel streaming must be open, so any data type or file format can be streamed. Synchronized playing of live or stored streams from any server to any client is also a requirement if streaming is to become ubiquitous. To this end they are working with many vendors, and Mr. Glaser listed many data types they now support.

Mr. Glaser summarized by stating that we are entering a new era of communications based computing, that key technologies and standards are aligning, early adopters are already deriving competitive advantage, and the time to deploy is now.

State of the Web: A Perspective

Ed Tittel, President LANWrights, Inc.

Ed Tittle is a well-known author and consultant and member of the Interop program committee, specializing in networking with an emphasis on the Web. His presentation was one man's predictions about Web technology, not tools or demographics. Natanya Pitts-Moultis collaborated on the preparation, but Mr. Tittle made the presentation.

He began with a disclaimer, stating that the past can predict nothing new, and that revolutionary changes can and do occur. For example, he asked how many people would have predicted the rise and fall of Active-X three years ago.

His talk covered four technologies:

1. HTML 4.0
2. Cascading Style Sheets (CSS)
3. Dynamic HTML
4. Extensible Markup Language (XML)

HTML 4.0

The HTML 4.0 recommendation came from the World Wide Web Consortium (W3C), on December 18, 1997. It comes in three flavors, one requiring strict adherence to the specification, one that is a transitional state between 3.2 and 4.0, and a separate one for those using frames.

It took W3C 26 months to get from version 2.0 to version 3.2, which gave vendors time to create and propagate proprietary features, and version 4.0 is an attempt to unite browsers. The separation of presentation from content with style sheets is a big step in this direction. Tables and frames are also improved. Unicode is also supported, allowing us to get away from English and the Roman alphabet.

His advice is to move to HTML 4 in order to ease support of multiple browsers and hardware configurations. Mr. Tittle predicted that HTML 4 will create a new lowest-common-denominator for Web content, but he still feels that it fails to offer truly customizable markup or dynamic behavior.

Cascading Style Sheets (CSS).

There are two CSS versions:

CSS1: Recommended, http://www.w3.org/TR/REC-CSS1

CSS2: Working Draft, http://www.w3.org/TR/WD-CSS2/

CSS1 fulfills the same function as templates in Microsoft Word or other word processors and page publishing programs. You can define styles by name, specifying

things like font size and types and text colors, and you can position items precisely on the page. You can establish a list of fonts to be used if the user's machine does not support your first choice.

There are also color controls. For example, you could set level-1 heads to be blue by default, yet override that at some point in your document. You can even assign behaviors to color, for example, having it change when the mouse moves over some text.

You can create a library of style sheets which can be referenced externally. This allows a central design group to maintain a library of styles to be used by the organization, providing control and uniformity of appearance, and frees people concerned with content from having to be make style decisions. Mr. Tittle recommends moving to CSS1. (Natanya Pitts-Moultis has written a book on CSS1 for Coriolis Press).

CSS2 is a strict superset of CSS1, so your CSS1 work will be protected. It supports different media types, which means you can provide alternative forms of delivery for the same content. For example, you could offer sound in place of video for blind people. (You may be legally obligated to do so).

There are also font handling extensions, like allowing a different font for the initial character in a paragraph and other text effects you are used to seeing on printed pages. In addition, there are style controls for tables, for table headings, cell data, padding, etc. Layers and dynamic positioning of elements on the page are provided and libraries of standard headers and footers are possible. You can even associate a different color palette with a page as a function of the user's browser.

CSS make it easy to enforce a consistent look and feel and separate content from presentation; however, one must learn a new, more complex markup and syntax. Mr. Tittle predicts that CSS be commonly used for complex sites, but may be overkill for simple ones.

Dynamic HTML.

Dynamic HTML allows dynamic client-side behavior, without requiring traffic between the client and server. Appearance may change and local updates may be made without involving the server. Microsoft and Netscape do this in somewhat different ways, and Microsoft has even copywritten the term "Dynamic HTML" (with a capital D). W3C is working on their specification, but the work is going slowly, and in the meantime, Microsoft and Netscape are charging ahead.

W3C is attempting to consider CSS, DHTML and XML together. This will enable things like changing the color of the third level-3 heading on the page after the document has been downloaded. This sort of thing can be done in response to user events like mouse clicks or moving the mouse over an object on the page.

Mr. Tittle suggests taking an experimental approach, but perhaps not converting your site until the standards shake out. The WC3 working draft can be found at:

http://www.w3.org/TR/WD-DOM/

WC3 has been more influenced by Microsoft than Netscape, partially because Netscape did some unpopular things like only supporting JavaScript as a scripting language, and supporting non-standard CSS.

Microsoft's Dynamic HTML has a full-fledged HTML object and event-handling model, and is compatible with CSS1. It provides a means for communication with a database on the server. There are also a number of multimedia controls, like canned dissolves and wipes between sequential screens. Voice and key commands can generally substitute for mouse commands.

Netscape has a different object model and style sheets, and dynamic behavior is achieved by displaying different layers. This requires re-rendering entire pages to make small changes. Netscape can only change content while the page is still downloading, unless you have multiple layers. This means, for example, that to play a long video file with Netscape, you would have to download it with the page; whereas, with Microsoft, you could download it only if the user clicked on the button to play it.

Dynamic, event-driven behavior is key to next-generation applications. This will move the Web from a document repository to a general purpose platform for many applications.

Mr. Tittel believes, however, that we are at least a year away from a standard. So, for now, you should either content yourself with experimentation only, or be resigned to having to deal with the

complexities of two platforms (unless you have an intranet and can control the user environment).

Extensible Markup Language (XML).

XML stands for eXtensible Markup Language. You can find the most recent recommendation:

http://www.w3.org/TR/1998/REC-xml-19980210

XML is a metalanguage, which means it is a notation that allows you to define custom markup languages. You must define tags with a rigorous document type definition (DTD). XML is a subset of the Standard Generalized Markup Language (SGML). HTML is an example of one tag set which may be defined via an SGML DTD.

Since you are free to define your own tags, XML provides great flexibility. There are over 35 known XML-based markup languages already in use. Examples are the Channel Definition Format used for automatic software maintenance, Chemical Markup Language for molecule description, and Mathematical Markup Language based on Donald Knuth's TeX.

XML has the potential to allow the definition of multimedia data types, eliminating the need for plug-ins. Via the Synchronized Multimedia Integration Language, sound, graphics, transition effects, etc. will be synchronized. There is also a Commerce Net initiative under way to define an XML-based markup language for electronic commerce.

XML will provide great flexibility in data types, linking, and application behavior, but it requires a working language of SGML, which is no small task. Charles Goldfarb, one of the inventors of SGML, is editor of an SGML book series at Prentice Hall, and this would be a good starting place for learning. Tittel predicts that over the next 5-10 years XML will make its way down from the high-end to general acceptance.

Certification and Commerce.

The Web is becoming a global marketplace, but electronic commerce has too many standards to settle down yet. The same is true for privacy via encryption, identity masking, cookie refusal, content selection and blocking, identity verification, authentication for content and code, restriction of what executable code (like Java applets) can do on your machine (you may wish to allow software from certain sources to have more rights than other software), and secure storage of sensitive data.

Solutions to these problems include certification and use of digital certificates, electronic wallets (*how do you synchronize your desktop and portable wallets*?), secure electronic transactions, the XML-based Commerce Net industry initiative mentioned above, eCash, and micro-payments.

Since there is so much money to be made in e-commerce, Mr. Tittel is confident that the standards and systems will eventually become available—it is inevitable. For now, either plan on spending at least $5,000-10,000 on a turnkey system or on spending a lot of time.

Mr. Tittle threw in a few additions to his prepared talk. He cited a Gartner Group study stating that it costs $15 a year to maintain a dynamic Web page and $225 a year to maintain a static web page. He mentioned that servelets—server-side applets that extend back-end functionality—are gaining importance and that people are now beginning to plug Web servers into document management systems, thereby cutting maintenance cost. (This has been done with databases for some time. It is the sort of thing that reduces maintenance costs from $225 to $15 per year).

For slides from this presentation or additional information, please see:

http://www.lanw.com/

Training and Distance Learning
Moderator: Al Gordon, Program Manager Siemens Virtual University

This session focused primarily on low bandwidth, instructor-directed training which combined synchronous and asynchronous material. The application is at Siemens and the server software from Databeam. There were three speakers, Yatman Lai and Al Gordon from Siemens and Debbie Black from Databeam.

Yatman Lai
Manager for Learning Technology
Siemens Virtual University

Mr. Lai gave an overview of Siemens view of the role of training and distance learning. Training is strategic at Siemens, and the training department is challenged by

globalization, competition, and demands to increase productivity while cutting cost. Siemens also has self-directed work teams, with no time for classroom training, so continuous learning is needed. Training is key to attracting and retaining people. Distant mentoring, training and collaboration are all important.

Education was traditionally done in a group learning environment with a captive audience. Computer-based training (CBT) like CD-ROMs allows any-time, any-place individual training, but you are isolated and on your own. They seek a middle ground with some material on-demand and some interactive:

On-Demand

- curriculum planning
- course catalog
- course schedule
- course enrollment
- Web-based learning

Interactive

- discussion groups
- data conferencing
- application sharing
- audio-graphic conferencing (and later video)

Today companies have intranets with information scattered around the organization. But ideally they want to have learning-related material available in one place on the intranet, at a "virtual university," with reference material and Web-based training.

This material must be augmented with live, interactive training—"the glue that holds it all together." The interactive material may be a dialog with someone, an expert you can ask a question of, discussion groups, data conferencing, audio-graphic conferencing, application sharing. Ultimately video conferencing will be widely used, but it is still too early.

Their goal is to make learning less intrusive (and time-consuming). They want to create a classroom on every desktop and laptop and a network of interactive learners. You do not ask *"what course can I take*?" Instead, you look at the whole knowledge repository of the company, and ask "***what can I learn***?"

They will support continuous learning, collaboration, mentoring and independent workgroups on the Net. Success factors include fostering a company culture that embraces networking, a tremendous amount of knowledge to be shared within the company, and a strong relationship between the training and IT departments (which is now strategic).

The line-of-business organizations are also important and must value their role in the transfer of knowledge. Employees must also see training as worthwhile and valuable. Finally you must promote it. Training and education—the management and dissemination of the organization's knowledge—is the intranet "*killer application*."

Debbie Black
Market Development
Databeam Corporation

Siemens uses Databeam server software for their virtual university, and Debbie Black talked about their view of distance learning and their products.

The technology for online collaboration has improved in recent years. It began with expensive video conference rooms for high level meetings, using proprietary audio and video formats in the 1980s. In the 1990s, desktop conferencing and standards began to emerge. In 1996-97, we saw the transition to IP networks using low-cost equipment with standards in place. The next challenge is cultural issues, and Ms. Black expects a relatively large share of cost and effort to be in content and consulting rather than hardware and software.

She gave the audience some tips:

- Training needs to be live and interactive at least at some times.
- It must also be Web-based and easy to access. Databeam now uses thin clients and Java applications so the user has transparent access to easily-used software.
- The interactive portion of training must also be well integrated with self-paced material.
- There must be an alternative for those who miss group meetings or other synchronous events. You can also integrate Internet material with your Web-based training.
- You must know your audience capability and how they connect (e. g, their speed of connection). It is critical that you adhere to standards: T-120 for multi-point data communications and H-323 audio and video communication. RSVP and other forthcoming service-quality guarantee protocols will

be compatible with these when they are deployed.

ISPs are also deploying servers so that you do not necessarily have to operate your own. For example, MCI offers a Net-based conferencing service, and will set up an audio or data conference for you on demand. The same may become the case for distance education.

There is a tradeoff between the amount of interactivity and management complexity. Browser-only operations are the lowest end. Next comes Java applications, but you must ensure that users have Java-enabled browsers, the application can get through firewalls, etc. Full application software like Microsoft NetMeeting is most powerful, but must be installed and capable of operating.

Databeam has two IP products, a conference server and a learning server. They are used in four general areas: brainstorming and planning, team meetings, large online events, and virtual classrooms. A company typically runs both, and Ms. Black gave examples of each from marketing, help desks, training, company meetings, etc.

Al Gordon
Program Manager
Siemens Virtual University

Like all companies, Siemens is confronted with fixed function jobs evolving into expanded roles, rapid technology change, short shelf-life of knowledge, and the need for continuous learning rather than events like classes. For the training department, this means more courses and instructors, but the reality is that training budgets and headcounts are not increasing.

They are under pressure to cut field expenses and to improve productivity. Siemens provides training and education for 5,800 employees in the US, and has begun working with distance learning.

They had good infrastructure, including a nation-wide LAN/WAN, and introduced an intranet with good remote access capability. They also have modern desktops with Pentium PCs running Windows 95. With that technology base, they launched their Virtual University for instructor-led synchronous education.

They decided today's low-bandwidth video quality was too poor, so they went with audio-graphic conferencing—separate, but simultaneous data and document conferencing—carried over IP networks and using the Databeam conferencing server. This is combined with a simultaneous conference call using the public-switched network (not telephony over IP).

They provide instructor-led synchronous education. Mr. Gordon cited data showing that less than 12 percent of purely asynchronous (self-led) CBT courses are completed in the corporate world. They combine synchronous and asynchronous training, with students doing CBT and taking tests between on-line class meetings.

The program was low-cost (under $30,000 hardware and software) and easy to implement relative to video conferencing. Education is at the desktop—office, home, a hotel, etc.). This is the precursor for desktop multimedia (which will be ready in three or four years).

They are using a medium sized (233 MHz Dell) server with moderate specifications, Windows NT, Microsoft's Internet Information (Web) Server , and Databeam's server software. They have their own audio bridge for the voice conference. (At first they used commercial telephone conferencing, but realized they could pay for their own bridge in about two-months billing).

They also have a remote access system. They initially used FarSite and NetMeeting clients, but found them cumbersome to install and not well suited to their applications, so they are moving to Java and thin clients. Users also began using the video for many other applications and for playing around, so they went to a stripped down version of NetMeeting that did not allow end-user conferencing.

The instructor and students can talk, share a whiteboard, use a highlighter or drawing pen, share applications, etc. They find this much more compelling than PowerPoint presentations.

Students can save presentations and material can be imported from common graphic formats and any Windows-based program like Excel, Word, or Power Point. They can also do screen and video capture.

Mr. Gordon presented an evaluation of their first course. The task was re-training over 500 telecommunication engineers. This had traditionally been done in an 8-day classroom program. They replaced that with eight 2-hour audio-graphic sessions, each

separated by a week for asynchronous assignments and online testing. Classes were 22 students and 2 instructors, and the material was advanced and at the college level. The topics included:

1. Standards: OSI, IEEE, and ITU
2. LAN/WAN: Client Server
3. TCP/IP
4. Internet, WWW and Intranet
5. ISDN: PRI, BRI, and Broadband
6. ATM, FDDI, and Frame Relay
7. Video
8. Wireless, Cellular, and PCS

They found the class was compressed—there was no fluff or padding which may occur in a fixed-length class in a classroom. Instructor platform time was reduced by 75 percent. They saved sending people to a training center for 8 days. (Interviews indicated that the workers worked extra hours on session days to make up the time lost rather than accomplishing less). The savings for each class of 22 students with two instructors were:

Savings	
Travel/Lodging $3,155 x 24	$75,720
Field Productivity $1,650 x 22	$36,300
Expense	
Telephone charges: 1,080 minutes x 24 x .265	$6,869

This results in a return on investment in hardware, software, development time and presentation time of over 50 percent. (The instructors re-purposed their classroom materials, and development costs are not included here).

They carefully analyzed anticipated loads before starting, so the training application did not swamp their network. Pedagogically, they found that frequent student interaction is important, people in their offices must be able to make the space and avoid unscheduled interruptions, and students liked the time between sessions to solidify knowledge.

Subsequent to this evaluation, they replaced Databeam's conferencing server with their learning server, and went to Java on a thin client. The learning server gives them an explicit participant list, live questions and answers, and student "hand raising." They have also implemented recording and playback so a lesson can be recorded by the user, and purchased their own audio bridge to save the telephone charges.

Multicast Grows Up

Moderator: Jeff Kunst, Newbridge Networks

This session featured three speakers, Martin Hall, Rod Murchison, and Mike Galli. They updated the audience on the state of the IP Multicasting, which they feel is ready to move out of research and development and be deployed on the Internet and intranets.

Martin Hall, Stardust Forums

Mr. Hall's company, Stardust Forums, is the organizer of the IP Multicast Initiative (IMI), an 18-month old group of users and vendors working on standards and multicast promulgation. He noted that audience sizes at his presentations have grown steadily since he began making these presentations when the Initiative was formed.

Mr. Hall discussed IP Multicast Background and technology and its current business and technical status. He also described the IMI.

Multicasting is useful if you have material—perhaps an audio or video stream—that you wish many people to receive simultaneously. It is the Internet Engineering Task Force (IETF) standard for multi-point communication. It was invented ten years ago by Steve Deering at Xerox Palo Alto Research Center, and is often used to cover events in the research community. (Deering is now at Cisco Systems.) The protocols, which matured within the research environment, and are now being implemented commercially.

Multicast is designed for efficiency in transmitting a message from one source server to multiple receiving clients:

Unicast	one sender, one receiver (the typical Internet mode)
Broadcast	one sender, everyone receives

Multicast	one sender, "tuned-in" group of receivers

The goal is never to send more than one copy of the same data across the network.

The client computer initiates the transaction that tells its local router to tune into a particular multicast stream. Clients can join and leave groups at will using the IGMP protocol. Therefore, to multicast, routers must support the Internet Group Management Protocol (IGMP) and multicast routing. Routers and switches that do not support these functions may be upgradable.

Multicast traffic is only routed to destination networks which have at least one client tuned into the particular multicast address, thereby reducing load on the server and network traffic. Where you may need a group or "farm" of servers to support a large number of simultaneous users in a unicast environment, one server may suffice in a multicast environment.

When the multicast initiative began 18 months ago, there was already a good deal of support in hardware—routers, switches, etc.—and it is also now supported in operating systems. Streaming audio and video application software like that of RealNetworks and Microsoft now also support multicasting.

Congestion and bandwidth scarcity are always issues on the Internet. If you have a T1 connection, you can only support 55 concurrent users at modem connections. With multicast, a single stream would fan out to reach all tuned clients.

The Internet is increasingly going to be driven by content. AOL, Netscape and others are becoming media companies. Reaching large audiences is key to advertising-supported services, so we need the efficiency of multicasting. (AOL has a larger audience today than MTV).

Multicasting offers a value-added service opportunity for ISPs. They can provide multicast service to customers, either on proprietary extranets or the Internet. UUNet is offering such a service now.

This is only the start of demand for multicast-compatible applications. As high-speed last-mile solutions like xDSL and cable modems become widely available, such applications will proliferate rapidly. High-speed wireless—whether terrestrial or satellite—will also facilitate these applications. Networks are at capacity today but new investment is occurring rapidly.

The IMI has around 80 member companies. These include leading hardware and software vendors, users, ISPs, and cable, satellite, and telecommunication companies. Its goal is to facilitate broader deployment of multicast on the Internet and within corporate intranets.

The IMI organizes technical programs and provides technical resources for the members. They are seeing much more deployment this year than in the past, so there are now a series of deployment programs and training sessions. They have materials such as a deployment guide (which is available on request). They also run an Internet server to keep help vendors and users find each other in the marketplace. From June 1-5, 1998, IMI will run the first of what they plan to be an ongoing series of interoperability tests.

Multicast has moved from research, to commercial products, and now to deployment. It is necessary if we are to support audio and video on the Net of the future. For more information, see:

http://www.stardust.com
http://www.ipmulitcast.com

Rod Murchison
Multicast Product Manager, Newbridge Networks

Murchison talked about deployment considerations and issues, and concluded with mention of capability we can look forward to in the future.

Before deploying IP Multicast, you need to study all aspects of your network and applications. For example, if many users need online training, multicast should be considered. You need support in your hosts and routers for multicast. If you have to upgrade hardware, it will be costly, so you should consider the cost/benefit carefully.

IGMP allows a client to tell its router and other hosts on the LAN which stream it is receiving, therefore all hosts and routers must support IGMP. There are three versions of IGMP

- In version 1, hosts left groups by not renewing membership, which meant high overhead.

- In version 2, one stays tuned until a termination command is executed, significantly cutting bandwidth requirements.

- IGMP version 3 is an IETF draft at present, and it will allow multiple senders and receivers, for example, for a multi-person conference.

One major concern is the broadcast of packets to all peer hosts on a LAN even if only one is tuned to a multicast stream. This may place a significant burden on the network interface cards and processors of hosts which are not receiving the stream. (They must decide to discard the extraneous packets). This will always be the case if you connect hosts to a hub or most switches.

Some modern switches, though, now do multicast containment, to eliminate this problem. The switch "listens" to IGMP messages, and acts appropriately. Your hardware may have to be replaced or upgraded to do this, and you should definitely look into this if you plan high bit-rate applications.

Murchison also summarized the four multicast routing protocols which are now in use, and the factors to consider in selecting among them. One factor to consider would be the percentage of clients on a network that frequently receive the same material. Network configuration is also important in choosing the appropriate protocol. If one server is a common source for most programming, you would tend to different a different protocol than if there were many. Protocol selection is an important deployment decision—you should consider the distribution of typical senders and receivers.

Multicast will work well with high-speed routers and ATM in the future, and selectable quality of service will be available. Encryption and authentication will also be available, so you can protect sensitive material and limit access. Directory-based usage monitoring and billing are also being worked on.

Mike Galli
Director of Marketing
Optivision

Optivision has traditionally worked on MPEG broadcast, but is now working on enterprise networks. Mr. Galli mentioned a number of things which remain to be done:

- Optivision sells video encoders which can generate multicast MPEG streams, but more development is still needed on servers and clients (the deployment of IGMP version 2).

- Analog broadcast users must learn new skills and become concerned with the sort of network design and protocol selection issues discussed by Murchison.

- User-friendly client software, which makes the computer look more like a TV with a VCR is needed.

- ISPs want relatively uniform traffic throughout the week, whereas multicast traffic peaks and drops off.

- Tools are needed for managing content and keeping track of where it is stored and in what format.

Mr. Galli assured the audience that work is underway in each of these areas.

Extranet War Stories

John Kay, Chrysler

As manager of electronic commerce for Chrysler, Mr. Kay is responsible for communications with suppliers, including EDI, email, the company's extranet, and any other technologies suppliers want to adopt. The extranet—www.spin.chrysler.com —can be accessed by the 12,000 suppliers that have a business relationship with Chrysler by using their password and user name.

Chrysler's EDI arrangement, which includes 17 transaction sets, is the best in the industry because the company has taken a proactive stance to get all suppliers online. Use of it has virtually eliminated most other forms of communication required to do business with Chrysler.

Although EDI has been successful, Chrysler does not have 100 percent participation, especially by smaller suppliers, because EDI is expensive. Thus, Chrysler is introducing EDI over the Web. It has 2,500 non-production suppliers online, it is trying for another 4,000 via the Web.

The company has 45,000 users on Lotus Notes, generating lots of online communication. Chrysler ran into a few problems with Notes. One was cultural—not copying others on submissions. Externally, an issue was unsecured communications. Furthermore, there is a limitation to the size of attachments. Finally, it would be desirable if industry email addresses were stored in one

directory. At the moment, they are not.

Chrysler's extranet started in 1993 and used proprietary software. More recently, the company moved to Web technology. Some 5,000 suppliers are now online, with 14,000 user IDs. The extranet contains 45 applications—a mere 10 percent of their goal. Ten of the applications are mainframe-based and the other 35 are Web-based, where the users ask questions or seek information.

Typical queries include, "When will I be paid?" "What information do you have on X." "What is your policy on Y?" And so on.

When initiating the extranet, Chrysler asked for input from various departments—legal, audit, patents, corporate communications, and IS. It was a bit tedious, but Mr. Kay discovered that these departments were only interested in setting policy at the outset. They had no interest in monitoring the site. Asking for their input at very beginning got his group off to a good relationship with these departments because it showed he was concerned with corporate policy. Then they let him build the site his way.

His e-commerce group approached management of the site construction by saying, "Here's a great opportunity to communicate to external users. Management bought the idea and provided the funding. However, over time, unless the business benefits are continually emphasized, they see it only as an IT project, not a business project. It was imperative to continually demonstrate all the business uses of the site." In fact, Kay's group even created presentations, cost/benefit analyses, and pamphlets to sway employees to get behind the technologies so they could communicate with suppliers on an equal basis.

Mr. Kay discovered that the best strategy for working with end users was to ask what they wanted online—rather than ask what they thought dealers or suppliers should have. By asking about their own needs, he got support for future applications.

As you understand your audience, said Mr. Kay, be aware of the hardware and software they have. Also realize that some users are techno savvy, while others are not. Keep your site simple—at the lowest common denominator—because you never know what knowledge or equipment the user has. For example, only support Web browsers after they are in production. You do not want your suppliers to be beta testers. Also, do not include Java applets that cannot go through a firewall.

Over the past four years, Chrysler has learned that the extranet needs to keep changing, to keep up with the fast-paced world of today. Their security was shoehorned into the extranet from the mainframe applications. The sign-up process was laborious, requiring some 15 pieces of data. They learned this was not tenable for Web users because many suppliers, governments, state senators, distributors, and others want a relationship with Chrysler. Registering them and keeping the security data up-to-date was ponderous. Chrysler is therefore tying to move security on to the suppliers to manage themselves.

In addition, Chrysler has learned that its extranet has numerous redundancies with Ford and GM. For instance, users must remember three of everything. An industry action group is therefore creating a common user ID so as to unburden end users. The industry group is also experimenting with certificates to authenticate and authorize users.

Watch traffic at your own site to see what's hot. Then try to figure out what else users might want to see, advises Mr. Kay. One goal at Chrysler is to find out what users do after visiting the first page. If you get 1000 hits/day on the home page, but 30K on other pages, your site is doing something right.

Chrysler's first extranet site lacked luster, so it was redesigned by categorizing applications by business area, for easier use. The left side of each screen lists the general areas, such as purchasing, sales and marketing, and so on.

Chrysler discovered that users misinterpret information they find on the site, so training is important. Training also pushes responsibility to end users, which is important. Chrysler's goal is to have more external than internal users by the end of 1998.

They also learned to have first-time user pages. These give an overview of each application and the user's responsibility for each one. In addition, the user manual is online, along with a contact list of application owners. But Chrysler encourages self-help, embedding hot links to help pages.

The Big Three automakers are realizing they need to help each other on the administrative aspects of managing their extranets. For instance, they are thinking that the sites should look somewhat similar, so that users only need to learn one interface. They plan to use the same

help keys, always have the same contact list in the lower left, and so on. They want to reduce the burden on users to interact with them.

The Big Three are also working on an industry-wide email directory and creating an electronic commerce yellow pages, which will contain the business requirements of conducting business online with each company—how they handle email, attachments, and so on..

Managing an extranet is like driving at 80 mph. Things come and go pretty quickly and you have to stay on the ball. Web strategy changes every 90 days. Furthermore, applications pop up quickly. If the infrastructure is not in place to put them up rapidly, you have problems. In addition, when you launch a new application, you will have users who access it that same day. You need to be ready. Before the Web, people launched applications without 100 percent quality. That does not work with an extranet.

Currently, Chrysler is concentrating on linking internal and external users, reducing complexity, finding more applications for external users, and adopting industry standards.

Dave Leonard, Corporate Express

Corporate Express, which was founded in 1986, is the world's largest supplier of non-production corporate supplies and services. It sells office supplies, furniture, basic software and computers, janitorial supplies, and such. It is commodity- and service-driven with sales of $4 billion, over 500 locations, and 80 distribution centers in 11 countries. It has 28,000 employees and 10,000 delivery vehicles. Technological change has helped sustain its growth and customer ease of access.

Corporate Express' technology objective is to improve the supply chain of the non-production goods industry, making it very efficient. Since this is a commodity business, competitors cannot differentiate themselves on products, so they differentiate themselves on service. The goal is to create customized solutions for every customer.

The company's role is to shield corporate customers from the complexities of dealing with many vendors of non-production goods. It is expensive for corporations to deal with 6,000 suppliers. It is much more efficient to deal with 15, or one. Corporate Express owns a large number of companies, so it can include all, or most, orders under one invoice. It also has relationships with other companies, to also ease a customer's ordering process. For one car company, Corporate Express was able to reduce the number of suppliers from 1,600 to 50.

Customer-facing technologies.

Corporate Express custom-built an electronic infrastructure to make the company more efficient. It uses this infrastructure to take orders, serve customers, plan deliveries, and control reporting and finance. The infrastructure has progressed through four technologies: EDI customer-premise catalogs, the Internet, and virtual private networks.

EDI.

If a customer has an EDI or storefront management or ERP system, Corporate Express can accept their applications in any EDI format. They have several Internet capture systems to capture orders. On the back end, they source order to suppliers, such as a company that sells hammers. In essence, they act as the procurement switch.

Once it has been set up, EDI is operationally very stable. But since Corporate Express tries to fill orders the same day for next day delivery, EDI does not meet its needs because it uses store-and-forward technology. In addition, Corporate Express owns all the problems with EDI.

So the company's extranet started with EDI via a value-added network. The company has 150 EDI users and is receiving over one million EDI-based orders a year.

Customer-premise catalogs.

The company then installed catalogs on customers' computers (AS/400s to PCs) and a network link to transmit orders. The company has 850 customer premise catalogs and is receiving 250,000 orders a year via this route. But this method of ordering is not increasing fast.

Catalogs are stable and customers love the functionality, but they are difficult to implement. You have to install them in the customer's site and they are not scalable. Data management is also difficult because updates must be done at each site, and catalog prices change often. As with EDI, Corporate Express must fix every problem, even when the problem is with the customer's NT server.

The Internet.

The company then created an extranet for Internet-based ordering. The company currently has 600 corporate customers ordering via the Internet, and it is receiving 300,000 Internet orders a year. But this

method of ordering is increasing quickly.

Internet ordering is a great success. Customers love the functionality, and data management is easy. However, it can be technically difficult to implement. It requires 24/7 availability and Corporate Express still owns *every* problem—even if the problem is, say, due to an outdated firewall at the customer's site.

Virtual private networks.

The future is Internet tunneling, extending the company's intranet to customers via encrypted and virtual private networks. The company will then be able to pump orders directly into its ordering system. But it has only experimented with tunneling with two customers and has received only two orders. The company has discovered that every time the proxy server was changed, they encountered problems. So the company still seems to own all the problems. But, this technology holds great promise.

Don Wilde, Soligen

Soligen is a $6 million company that makes functional metal castings in days (rather than weeks) for almost any size or metal alloy. They make complex metal parts for auto suppliers and tire companies. Their niche is rapid manufacturing—supplying faster parts or better parts at lower total production costs. They accept a 3D CAD file and return a finished part in two weeks, even if it is very complex. Engineers use this service to quickly get a part for testing. Then they redesign it, order a new version and test it. And so it goes through many iterations.

The company started six years ago in a garage where six engineers designed a machine to build ceramic molds layer by layer, fire them, then take them to a foundry to make a sand casting mold from which the parts could be cast.

They took their proprietary technology and broadcast it through the Internet to connect with suppliers and customers. They have thus created an extranet for getting requests and giving quotes, on a worldwide basis. The customer portions of the extranet are password protected. The company also has part of the site for investors.

The company receives orders from companies where the engineers do not speak English. They only know the correct numbers to plug into Soligen's site.

Soligen learned that when you deal with large companies, you must "be there" to talk to them. The goals of the extranet were therefore to further Soligen's business with Fortune 1000 companies, to allow more people to do business with Soligen, and to integrate their electronic CAD with customer design centers in other parts of the world.

Mr. Wilde believes that the Internet has everything now. He can have one corporate site and numerous design sites, all of which interact over the Internet, giving him a global presence.

Building the extranet with freeware.

Soligen created an internal ordering system in 1995. Customers did not want not use it though because they did not want to load his system onto their computers. So Mr. Wilde pulled back and went to the Web.

The first generation Web site was designed by an external contractor. It was introduced at the 1995 Autofact show. Then Mr. Wilde became the company's full-time IS department. The second-generation site had three times the functionality. It not only allowed customers talk to Soligen, but they could request quotes, and receive quotes back through the Internet, with email and fax confirmation.

The system has grown from PC based to Berkeley UNIX—which is free. Mr. Wilde uses stable release 4 of Internet BSD UNIX. Everything on the site is based on freeware—the Apache Web server, the network drivers, the Web browsers, the graphics manipulation tools, and so on.

"There is nothing wrong with freeware," says Mr. Wilde, "so don't be afraid of it." One-half the servers on the Internet use Apache, incidentally. Yahoo is based on BSD. And the entire DNS system that routes connections is freeware. So everyone is already dependent on it when they use the Internet.

Mr. Wilde had to have a system that did not crash because all his production work runs through the site. So he prefers freeware, because, he says, Berkeley UNIX is very stable. It does not crash. Windows NT and Novell cost more, they require more engineers, and they crash.

Since they were a small company, they did not want to invest capital in the extranet. So they created it by spending only $150,000, which went mainly for the contractors who built the first generation site and the two

programmers who worked on it part-time.

Freeware allows Mr. Wilde to work on such a lean budget because the Internet is his help desk. Whenever he hits a glitch, he broadcasts it on the Net and one of the original developers of the software in question generally responds within several hours with fixes. "Contrast this with trying to get something out of a major software company," he says. When you deal with corporate proprietary solutions, you will not get it fixed if it is not within their limits. With freeware on the Internet, the opposite is true.

Since Soligen's customers are large companies, it must work their way. Much freeware works and it costs less to build and maintain. But Mr. Wilde has to spend time figuring out the differences between the proprietary systems these customers use and the freeware he uses. Since he has source code for all the packages, this is easier. Furthermore, the Internet-based developers share their answers.

Mr. Wilde builds to industry standards, not to the latest buzz product. JavaScript is a problem, so he does not use it. New browsers are a problem, so he does not use them. Not everyone has Communicator 4.05. Many people stay with what they have because it works. Mr. Wilde also does not use special technologies. He stays with HTML 3.2 and his server generates pages dynamically via Perl. As a result, he does not have problems. Everything works.

Mr. Wilde believes that freeware allows taking a small company to being a larger company.

ISPs as Enterprise Partners for VPNs and Remote Access

Brian Barton, Cabletron Systems

Virtual Private Networks (VPNs) are at the forefront of people's thinking these days. A major question is: What role will ISPs have? These four presentations deal with:

- The underlying technology, standards, and protocols
- The hardware/software for deploying a VPN
- Information management for VPNs
- Going forward with VPNs

Rob Redford, Cisco Systems

An IP VPN is a way of getting private IP over a public infrastructure. Getting IP to your site may occur over the Internet, over a private network, or a private IP infrastructure. Once you get it, you can extend your intranet to encompass remote users—called a WAN intranet—or customers/suppliers, also called an extranet.

Three fundamental requirements need to be met:

- *Privacy*. You need a way to communicate privately with the your customers.
- *Predictable performance*. Some applications, such as email and store-and-forward application, do not require specific performance characteristics. But others, like video conferencing, do. The goal is for IP to recognize an application and classify it properly.
- *Policy*. A policy-based network means the ability to distinguish the different types of traffic on the network (such as the importance of each type). A policy designates certain traffic as high priority and other as lower priority. The goal is to have the network recognize this and transmit it appropriately.

Dial-based VPNs

The most common IP VPN is dial-based. These are mainly for single-user temporal-based connections. The two ways to build IP VPNs are via client-based tunnels or network access server-based (NAS) tunnels. The client-based tunnel starts on the client's PC. It dials into the network and creates a tunnel to the company intranet. This tunnel is built on the PC. A NAS-initiated tunnel is built by the service provider rather than the client. In identifying himself in a certain way to the ISP, the ISP realizes it needs to create the tunnel for that client.

Client-based VPN.

Here, a client runs a piece of software, such as IPSec, which tunnels through the Internet (or any public infrastructure), terminating on a firewall. Privacy is maintained because everyone creates his or her own tunnel. IPSec is an open standard that establishes the parameters that let two sites communicate securely. It uses such mechanisms as encryption, digital certificates, and device authentication.

Four of the main issues are:

- *Header information*. This describes everything else going on within the encapsulation. Two methods of specifying header information are via an authentication header or an encapsulating security payloader (ESP). These are essentially the same—they make the communication private—but ESP also hides who you are. Another person cannot tell the endpoints of the communication.
- *Key exchange*. IPSec is based on a public key encryption mechanism, so if you have someone else's public key, you can send them data, and only they can open the message using their private key. The problem is that you have to know who the other person *really* is. Key exchange uses digital certificates to exchange public keys, and then generates a unique session key for just your session.
- *Modes*. There are several mechanisms for encrypting the data, such as DES, Triple DES, amd RC4. IPSec does not care which you choose. It simply negotiates between the two sites.

NAS-initiated VPN

A NAS-initiated VPN—created by the ISP on the network access server—uses two basic protocols: L2F (Layer 2 Forwarding) and L2TP (Layer 2 Tunneling Protocol). Cisco invented L2F; Microsoft invented PPTP. Both do the same things, so the indsutry got together and created L2TP, a combination of both.

When you dial in, you are authenticated at the user site, the tunnel is then created on the server. There are two important points about this structure. One is that the security server is on the corporate intranet, so the enterprise can maintain all this information, keeping control of it, which they may wish to do. The other point is that the remote user is only being identified (perhaps based on domain, calling number, or called number), not authenticated. Authentication takes place afterward.

Value add is the reason to take this route. One advantage is that the enterprise does not need to manage all the IPSec software on all the clients. Also, since the service provider is creating the tunnel, they have visibility into the tunnel and can provide more value added services, such as premium services.

For example, they can route you to a premium tunnel, or make sure a certain number of dial ports are available. Or they could restrict the number of dial ports available to you and offer them as a very inexpensive dial-in service. They can also support private addresses.

Dedicated VPN connections.

These connections are permanently on and are typically for multiple users, such as remote offices or supplier firms. It gives higher performance and capacity.

There are three ways to build these dedicated connections: with IP tunnels, virtual circuits, and VPN-aware networks.

IP tunnels. An IP tunnel traverses an IP network using a security appliance or a router using a protocol. Dedicated networks can be build using IPSec, creating permanent tunnels rather than dial-up ones.

Virtual circuits. The second method to build a VPN is to create a virtual circuit network and then put a router underneath a frame relay or ATM connection. Although most people do not think of this as a VPN, it meets the definition. The public portion is the frame relay or ATM connection; the private portion is the router, which the ISP can then manage it. Privacy is being maintained by the virtual circuits under the router.

Generic Router Encapsulation (GRE) tunnels let you put IP packets inside another IP packet to send it across a network. This allows you to hide address spaces and hide routing protocols. In fact, an GRE tunnel looks like a point-to-point connection, and therefore get privacy. This is the standard way an ISP would build a VPN because it is easy to do with an IP infrastructure.

With a frame relay or cell relay network, an ISP can create a VPN by putting a router at the end of the connections, managing the routers and supply IP to the user. The advantage is that the ISP does not need to encrypt to create privacy.

VPN-aware networks. A new way to build them is VPN-aware private networks. Here, the network itself is aware of VPNs and manipulates them as part of the network fabric itself. This is important for an ISP.

When you are tunneling, the ISP cannot see the VPN. But an ISP needs to be able to see a VPN in order to add value to it. So this is the distinction in this new type of network. There is a lot of talk about

QoS. The only reason you need QoS is to implement your policies. Different data has different value and you want the ISP to handle it that way.

Getting differentiated services requires being able to have them delivered privately. This is what ISPs are aiming to do. A VPN-aware network is a connectionless network, like your own private Internet. it is not a point-to-point network. It also works much the same as a WAN, so it is easy to migrate to.

To really be efficient, the VPN has to built into the infrastructure so that all of the services can be delivered privately in a scalable way.

For more information, please see: http://www.cisco.com/vpn

Kurt Bauer, Ascend Communications

Deploying a VPN over the Internet: the hardware and software view. There are three types of VPN:

- Enterprise tunnels (either clear channel of secure channels)
- Frame relay and ATM VPNs (not IP but uses VPN concepts within the service provider infrastructure to deliver high value
- IP-based VPNs, a super set of Internet-based VPNs, designed to embrace frame relay and ATM to leverage the public and private IP sectors in the market today, as well as create private routing domains and delivery quality of service.

The first tunneling protocol was ATMP (Ascend Tunnel Management Protocol). It is a layer 3 protocol, which is a GRE-type protocol so it can deliver IP and IPX. With ATMP there is the concept of a foreign agent—the node with dedicated bandwidth from a regional site or is the dial-in termination. A home agent can reside in an enterprise, and be under their control.

ATMP has been deployed in two ways. One is router mode, where once a connection has been initiated from a regional office—dial or dedicated—the session is terminated and translated into a set of IP packets, creating connectivity to the home agent. This is how they establish simple VPNs via GRE, with inherent security. Typically, router mode is used by enterprises, for intranet/Internet/extranet applications.

The second option is gateway mode, which has a slightly different topology. Here, the home agent is moved from the enterprise into the ISP infrastructure, so they are controlling the home agent. The home agent can thus switch the traffic over frame relay or X.25 dial leased lines. The gateway mode is used primarily by ISPs to deliver multi-protocol VPN types of services to enterprise clients or for porthole sailing, to create very large subordinate networks from which other providers can buy services.

While all protocols seem to be converging into the L2TP protocol, it is important to understand how it compares with ATMP. The component of a network access switch (NAS) are LNS and LAC may be placed in different places. If NAS is within the service infrastructure. When the call comes in, whereas ATMP terminates the session, L2TP encapsulates the session to forward it through the network. A tunnel is created.

When the LNS is placed in the service infrastructure, the ISP can manage all the addressing. If it is in the enterprise, the customer dictates the addressing to be used. These are all important considerations for network design.

The clear advantages of L2TP is the consolidation toward it from all vendors because it allows multi-vendor interoperability. Encapsulation allows supporting multiple protocols (although it adds overhead). There are also inherent congestion control mechanisms built in, and there is tunnel authorization not just user authentication. It supports Network Address Translation, and can add gateway mode possibilities, making it more flexible. Most importantly, there is built-in extension to support IPSec.

IPSec deployments so far have concentrated on the enterprise, creating secure tunnel environments originating within enterprises. However, there is huge growth opportunities for enterprises to lean on ISPs to provide secure services.

But several issues need to be resolved first. One is multi-vendor interoperability. In IPSec, this is not as common as it will be within six months. Progress is being made. Another issue is dynamic key management. Many multi-vnedor implementations of IPSec orient toward static keys, which do not scale well. The third issue is the number of CPU cycles consumed in creating secure tunnels using dynamic key management. A great deal of scaling needs to happen, to scale up from enterprises to ISPs.

ISPs are expanding their Internet VPN offerings. Most of

them are happening outside the U.S. today. But it is a short hop to retread corporate networks (as companies did from X.25 to X.75) to make them Internet-based. It is well understood outside the U.S. and there is rapid growth within the U.S.

Most VPN implementations are aimed at offering value-added services. So they are being used for outsourcing and systems integration opportunities. There are many success stories about broadening services in this arena. It is a natural leap to get into new business areas.

Pilot projects are the norm for enterprise work in this area; although a few are "going for broke" and assuming their networks will be built on VPN technologies. The future is not wholesale changeover to ISP-run networks, but rather hybrid VPNs in certain regions, with specialized services offered. The hybrid approach is the most prudent approach.

David Kamm, Cabletron Systems

Information management for VPNs. The relevant issues: are network management for dial VPNs, LAN-to-LAN VPNs, and extranets (which actually may be a subset of the other two).

The following information is needed by a service provider depending on which VPN model they offer:

- *Phone book services*, especially for dial VPN. Enterprises have thousands of users, who have point-of-access numbers that need to be maintained and distributed via phone books. Much of this is in place today, but as dial VPN becomes more prevalent, it will be a larger hurdle.
- *Connection and tunnel statistics* for dial VPN. These are needed by ISPs to do adequate billing. This data is important if their revenues are based on who is connecting and for how long. ISPs might possibly also need authentication information; a subset of it might be held by the service provider.
- *Metrics to support service level agreement audits*. Customers may question service. ISPs need to therefore collect connection success rates for VPNs, connection bandwidth, and other such data.
- *Management of encryption parameters*. This is especially true if clients trust you with encryption key management, as well as set-up and tear-down of encrypted tunnels.
- *Cross ISP charging*, especially in the international coverage, where there could be settlements between ISPs. This is more prevalent in dial-up VPNs.

The following information is needed by enterprises for managing their VPN connections:

- *End user policies, authentication information and encryption parameters*. Most enterprises are going to want to retain control over this information. If ISPs have indeed offloaded this work, then they can charge to take it over.
- *Update and distribution of phone book style information* for dial VPN. This could be a shared responsibility, or totally offloaded to a service provider, at least for management.
- *Connection and tunnel statistics for both direct dial and tunneled connections*. Hybrid remote access is occurring, where traditional dialed remote access with some collected accounting data is being migrated to a VPN environment where there are logical PPP sessions. It would be nice to manage this seamlessly and with a single user interface.
- *Accounting information for departmental chargebacks*.
- *SLA metrics*, to ensure that the enterprise is getting the services contracted for.

If you compare what ISPs and enterprises need to know about dial VPNs vs. LAN-to-LAN VPNs, dial VP is more complicated. The main issues deal with whether companies are worried about authentication, initiating the tunnel, seeing tunnels, collecting tunnel statistics, and encryption.

There are a number of kinds of equipment used to house information about networks. These include

- *Connection management systems*—They all have separate pieces of the puzzle, and because they are isolated from one another, it is difficult to get an overall view. Most need extensions to handle tunnel information.

- *Directory servers*—In the remote access environment, enterprises want integration into the back-end systems for authentication and policy work in the dial VPN arena.
- *Firewalls*—of the firewalls augmented with encryption and tunneling, the Checkpoint product is the most visible today.
- *Client-end systems*—Windows (used by 90 percent of end users) needs VPN extensions, most notably the BaseCamp initiative. There are also third-party VPN products.
- *Dedicated VPN devices*—These are the most recent entrees. They combine all of the above, integrating firewall, routing, remote access, etc.

Integrating traditional and dial VPN access information into cohesive management systems is a challenge. **People are interested in migrating to VPN technologies. They want an evolutionary approach, not a revolutionary approach.** So we are hoping to see a common management environment to see dedicated WAN management along with tunneled (dial-up or dedicated) VPN WAN management.

We will see monitoring of SLAs, tracking what metrics are available. Not all SLAs will have readily available metrics and raw data.

Tunneling adds complexity to troubleshooting, because there is just a logical connection not the physical connection as in most PPP connections. This could lead to more fingerpointing among service providers. Expect to see traditional network management systems—such as OpenView—be enhanced with VPN management capabilities. There will also be shared management devices which give the enterprise manager visibility into some parts and the ISP manager visibility into other parts.

We will see MIBs in support of tunneling and VPNs. In fact, there is a draft L2TV MIB in existence. The sooner this happens, the better, because you can import those MIBs into standard NMS tools and have visibility and accessibility into multi-vendor VPN environments. Hardware vendors and VPN vendors are going to partner to have offerings.

Information management in a VPN environment is new and challenging, for ISP providers as well as enterprise network managers. Today's information repositories need extensions and better integration among them. We will see this as VPNs merge into mainstream. It is a fertile environment for ISPs.

If you can manage the VPN environment, you will have a leg up on competitors. In fact, you could even manage your customers' infrastructures.

Joyce Hervey, 3Com Corporation

Success factors in taking advantage of all this technology. ISPs first need to ask themselves what kinds of classes they will offer. Enterprise customers want to be able to prioritize the level of service they will get for each kind of traffic via policy management. Thus, service level guarantees need to take advantage of policy management and prioritize kinds of traffic. Service guarantees need to go beyond network availability to provide value added.

Enterprise customers are looking for network management tools that act as if they were operating the network themselves. They want their own workstation and they want to monitor what you are doing for them. They want live traffic data, so that they can see the number of users, the tunnels, and the amount of traffic throughput, so that they can help manage the SLAs themselves. They add value themselves by demonstrating that they are actually managing the vendors. So a CIO, when outsourcing their network, wants to know what kinds of hooks are in the tools and how those can be managed. The successful ISPs will have these.

Enterprise customers also want a valuable reporting structure that hooks into the service level agreement and guarantee. This is the same thing that happened with frame relay eight years ago. ISPs need to include this information to be successful.

Enterprise customers want scalability and shared control in the security arena. They do not want to endure a lengthy period of time to disconnect an employee when they leave the company. They want to know how soon they can turn off a connection. A huge value add is allowing an enterprise customer to terminate a specific user's ID, and other appropriate access, the moment that employee leaves the company. Enterprise customers also want the ability to create a demilitarized zone, as well as do quick adds and changes.

Billing settlement factors has become a sticky area. ISPs will have great success if they create either VPN billing settlement agreements

or partnership agreements with other ISPs—and ensure that those partners give the same level of guarantee as you.

Enterprise customers need to make sure that when providing their users with worldwide coverage, and for extranets and electronic commerce, that the network service providers have agreements in place. Customer want to be able to initiate a call anywhere in the world and know that the company applications will be available. This is the same thing that happened with X.25 with X.75 gateways, or more recently, with cellular phones—where roaming agreements were established so the users could call from anywhere. Prices may be outrageous, but the agreements do exist.

VPNs appear complicated to a lot of enterprise customers. They want to know how to tie in all the new kinds of equipment. ISPs are in a good competitive postion if they can quickly arrange a bundled offering that not only includes the customer premise equipment but also system integration and network management. To outsource, enterprise users want a comprehensive offering.

When going to a VPN, it turns out that some applications need to be changed at the customer premises because there are changes in such things as latency. If network service providers (NSP) can offer systems integration and value added services, they have a huge value proposition to offer customers. And it is an extremely high-margin business. Enterprise customers need this kind of help to get their systems installed because they have downsized their organization.

Industry success factors include using standard space platforms and interoperability among vendors. To help propel the technology forward, there is a move in the industry to have a VPN forum comprised of technology and enterprise leaders, not unlike what happened in the frame relay arena.

To be successful, enterprises need to understand a technology as well as their applications and the benefits of the technology. Once they understand these, they are in an attractive position to ask the appropriate questions of vendors. Enterprise staff will have to ask detailed questions of ISPs about QoS, hooks in their contracts, classes of services, and volume discounts for other kinds of communications. And they need to understand their corporation's level of risk and trust because outsourcing terrifies people. They also need to assess realistic expectations of what the vendors can actually achieve. To do this, trials using less critical applications are essential.

Editorial Analysis

If there is a single theme that ran through John Kay's talk about Chrysler's extensive e-commerce/extranet effort (see page 10), it would probably be success through careful planning and consistent execution.

As any developer will tell you, building a wired presence is easy. Building an effective presence that *successfully* links up thousands of businesses and tens of thousands of individual users is something entirely different.

Kay's talk reads like how-to manual on e-commerce. From the very beginning, the Chrysler system was built from a user, not developer, standpoint. This is crucial. All of the R&D in the world won't help you if your users are unable make business-critical applications work. Go right to the source. Your customers can be a provide a wealth of information on ways to improve the way you do business online.

His advice on design is right on as well. Keep it simple. A clean, straight-forward interface with common navigation is far more effective that flashy graphics and un-firewall-friendly Java applets. In addition, extensive tutorials help get users up to speed quickly. And once up to speed, it makes your job that much easier when you launch new features and applications because the users already know what to expect.

He also noted that the Big Three automakers are collaborating on ways to make their respective extranets consistent with one another. While this might seem to fall into the strange bedfellows category, it makes sense with Kay's user-focused development. If these three companies all use the same suppliers, a consistent look and feel across the various extranets would make their users' online experiences even easier. And for the less technically-inclined, this is critical.

This just underscores his more important point, that you absolutely must strive nothing less than perfection in all of your applications. Thanks to software patches and automatic updates, some vendors think that launching half-baked apps is okay. Perhaps

they are right—or not. But on the extranet, business-critical apps cannot have the same kind of touchy-feely period to work out all of the kinks. Just as they wouldn't release a car with an engine that didn't function, they cannot unleash an application on thousands of associates that wasn't up to snuff.

Interesting idea: build a system that is user-friendly, comprehesive, and actually works. As a user, that sounds good to me. In fact, I think you would be hard-pressed to find any computer user who was against something that would make his/her life easier. Perhaps some of software vendors could learn something from Chrysler. Just a thought. (*Contributed by Eric Rich*)

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
125 Main Street, Suite 475
Westport CT 06880
Tel: 203.221.1199
Fax: 203.221.1252

ILLINOIS

Giga Information Group
1300 West Higgins Rd.
Suite 116
Park Ridge, IL 60068
Tel: 847.823.2393
Fax: 847.823.2394

FRANCE

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

GERMANY

Giga Information Group GmbH
Carl-Zeiss-Ring 19-21
85737 Ismaning bei Meunchen
Germany
Tel (49) 89 9607 8302
Fax (49) 89 9607 8330

Giga Information Group GmbH
Konigsallee 60F
D-40212 Deusseldorf
Germany
Tel (49) 211.8903.272
Fax (49) 211.8903.288

ITALY

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

NETHERLANDS

Giga Information Group
Transpolis Schiphol Airport
Polaris Avenue 53
2132 JH Hoofddorp
The Netherlands
Tel (31) 23 568 59 52
Fax (61) 23 568 57 01

UNITED KINGDOM

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

Giga Information Group
Castle Hill House
12 Castle Hill
Windsor, Berkshire SL4 1PD
Tel: (44) 1753 831721
Fax: (44) 1753 831999

Subscription Order Form

Please sign me up for a subscription to the Conference Analysis Newsletter. I understand that I will receive a minimum of 20 issues per year.

❑ **1 Year**—$497.00 U.S./$552.00 elsewhere ❑ **2 Years**—$749.00 U.S./$849.00 elsewhere

Name: _____ Title:_____

Company: _____ Address: _____

City: _____ State: _____ Zip:_____

Phone: _____ Fax: _____

Please charge my: ❑ Master Card ❑ Visa ❑ Diners Club ❑American Express

#_____ _____Expiration Date (Month/Year)

_____Signature

❑ Check enclosed Bill My Company - P.O.#_____

Please make checks payable to: Giga Information Group.

Mail to: Giga Information Group, One Longwater Circle, Norwell MA 02061, or call: 800.874.9980, FAX: 781.982.1724. All checks must be drawn on a U.S. bank. Please add applicable state sales tax.