Section IV:

Multimedia, the Internet, and On-Line Services

High-End Digital Video Applications

Larry Amiot
Electronic and Computing Technologies Division
Argonne National Laboratory

The emphasis of this paper is on the high-end applications that are driving digital video. The research with which I am involved at Argonne National Laboratory is not done on digital video per se, but rather on how the research applications at the laboratory drive its requirements for digital video. The paper will define what digital video is, what some of its components are, and then discuss a few applications that are driving the development of these components. The focus will be on what digital video means to individuals in the research and education community.

The Digital Video Environment

In 1996, a group of people from several universities in the Midwest and from Argonne formed a Video Working Group. This body tried to define the areas of digital video of importance to their institutions.

It identified the following items:

- room-based videoteleconferencing systems
- desktop video
- broadcast video
- digital video servers
- virtual reality CAVEs

A virtual reality CAVE is a room that is roughly 15 X 15 X 15 ft. It has three projectors that are driven by computers which project images and actually generate a three-dimensional representation of a physical or mathematical model. By using special glasses, people can feel that whole images are floating around them, and that they can move around in them. That is one definition of virtual reality.

Video applications identified by the video working group include distance learning, research, collaboration, training, entertainment and information broadcasting, and workgroup conferencing. This paper will concentrate on those applications that most directly affect the research and education environment.

Video Infrastructure

One of the issues the Working Group examined was the communication infrastructure that drives digital video. Dedicated data circuits are being used in the infrastructure today along with integrated services digital network (ISDN)–switched circuits. High-speed dedicated circuits have the ability to transport high-quality, full-motion video. ISDN speeds limit the quality and motion that can be supported.

Internet and Intranet

The packet video networks which currently support many applications such as file transfer, Mbone video (talking heads), and World Wide Web browsing are limiting for high-quality video because of the low throughput one can achieve via the Internet or intranets. Examples of national packet switched networks developed in the last several years include the National Science Foundation Network (NSFNet). The Department of Energy had its own network called ESNET, and the National Aeronautics and Space Administration (NASA) had a network as well. Recently, the NSFNet was decommissioned, and commercial interests are now starting to fill that void. Research and education communities are finding, however, that this new commercial Internet is too restricting and does not meet their throughput requirements; it is starting to become congested. As a result, these organizations are moving back toward the intranet concept, and there probably will be another research and education network before too long.

Of course, a local site involves local area networks (LANs) utilizing technologies such as the asynchronous transfer mode (ATM), Ethernet, and the fiber distributed data interface (FDDI). Asynchronous transfer mode (ATM) is the one network technology that is applicable not only to the LAN (including Ethernet and FDDI), but also to metropolitan area networks (MANs) and wide area networks (WANs). It has the promise of providing not only high throughput but guaranteed service as well. ATM applications will be discussed in more detail below.

Examples of Digital Video-Driver Applications

A Virtual Field Trip

The first video driver application which I would like to discuss is a virtual field trip. In the fall of 1995, there was a conference was held in San Diego called Supercomputer 95. This was the big conference for researchers in the field of supercomputing. The organizers wanted to have the best supercomputers available for display on the floor, and they also wanted to link up existing supercomputers across the country on a single network called the I-WAY. In order to do that, conference organizers had to secure the cooperation of the various telecommunications vendors and use other existing research networks. They wanted a very high-speed ATM network, and they wanted to make sure that the network would be ready in time for Supercomputer 95. In May of 1995, a short, one-day test was carried out to see if ATM interoperability between vendors and carriers could be achieved.

FIGURE 1
High-Level Application Diagram–Schools & Applications (All ATM)

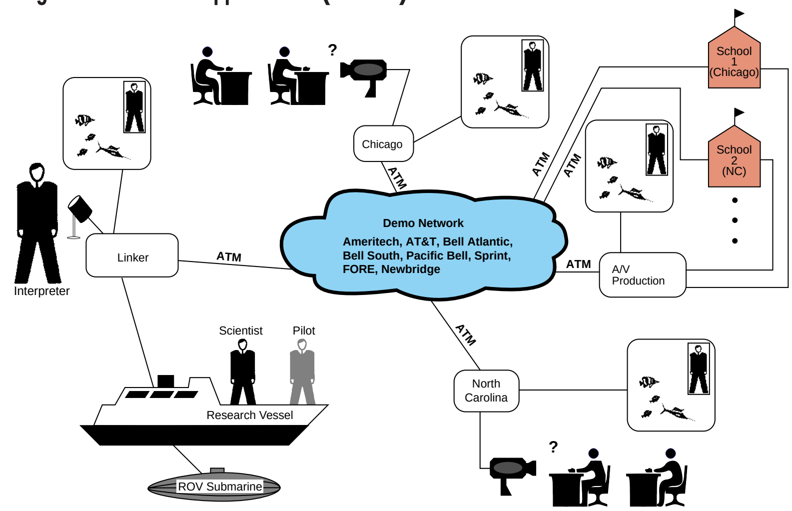

The virtual field trip took place in Monterey Bay, California, at the Monterey Bay Research Institute. The institute has a research boat that drags a submarine around Monterey Bay, and the submarine takes pictures of the fish, plants, and whatever else is at the bottom of Monterey Bay. This signal is then sent back to the aquarium, where an oceanographer can show the pictures to people sitting in a large auditorium and can lecture about them. The speakers has a huge digital server available, so they can put up a picture of something that is on the server's file system and answer the audience's questions.

Involved in the planning for the nation's newly emerging national information infrastructure plans for Supercomputer 95 was a committee in Washington, D.C. (chaired by the late Ron Brown, who was Secretary of the Department of Commerce) called the Information Infrastructure Task Force (IITF). It was part of Vice President Al Gore's National Information Superhighway Task Force. IITF chose to hold its first open meeting in May 1995 and invited the public and the press to see what was going on.

The organizers of Supercomputer 95 decided to transmit the pictures from the research institute in Monterey across a wide area network, all via ATM digital signal level 3 (DS-3) at 44.736 Mbps. The network used resources donated by a number of interexchange and intra-LATA carriers. The demonstrations were held in the auditorium at the Department of Commerce. Further, organizers wanted to get a couple of schools involved. They picked Pender High School in North Carolina and Whitney Young High School in Chicago.

Videoteleconferencing from each of these sites was set up using nationwide ATM links (see *Figure 1*)—from each of the schools, from the research institute, and from Washington, D.C. All this was done in two weeks' time. Whitney Young High School did not even have fiber-optic cable available, but Ameritech installed it within two weeks. Optical carrier level 3 (OC3), 155.52-Mbps fiber-optic cable was used to get the system operational. Essentially, the research institute showed some pictures and then switched to one of the schools, in which a classroom was set up where students would ask questions. All of this could be seen in Washington, D.C. Then, the oceanographer would answer the students' questions. At this point the picture would cut to Washington, D.C., where one of the secretaries might ask a question, and so forth.

The actual hardware that was put in place constituted a multivendor environment (see *Figure 2*). Different vendors' ATM switches were used. Newbridge was one of the vendors whose video board actually went in the Newbridge switch. Some of the other video equipment used at Monterey Bay included Sun Sparc workstations with parallax boards that fit right into the workstation. One board went into the workstation, and the other went into the ATM switch. Nevertheless, all the boards transmitted via ATM, and the images were sent across the country. Even though everybody used motion Joint Photographic Experts Group (JPEG) compression, the streams were all parallel. Video switching was done after the signals were converted to analog in the Washington location. There was switching at the analog end. In summary, this is what was accomplished in two weeks as a precursor to the I-WAY project. This is the type of application that is driving digital video.

Metropolitan Research and Education Network (MREN) Broadcast Video

In the Chicago area, there is the Metropolitan Research and Education Network, a network called MREN. It includes network systems at the University of Chicago, Fermi National Accelerator Lab, Argonne National Accelerator Lab, and the University of Illinois at Chicago. Ameritech is one of the vendor partners, and they have the central switch, which is located in downtown Chicago. It is a GCNS2000 ATM switch. Most of the links are synchronous optical network (SONET) OC-3 links running at 155 Mbps. There are still some DS-3 45 Mbps links in place, but very soon it will all be OC-3 with some people even considering SONET optical carrier 622.08 Mbps (OC-12). Northwestern University will soon join the network. This makes up a very high-speed network set up to help with research and education applications (see *Figure 3*). In addition to carrying just simple Internet traffic, video can also be put on this network (see *Figure 4*). Sun workstations have been set up, both at Argonne and at the University of

FIGURE 2

HTF Network Diagram ATM Demo Portion

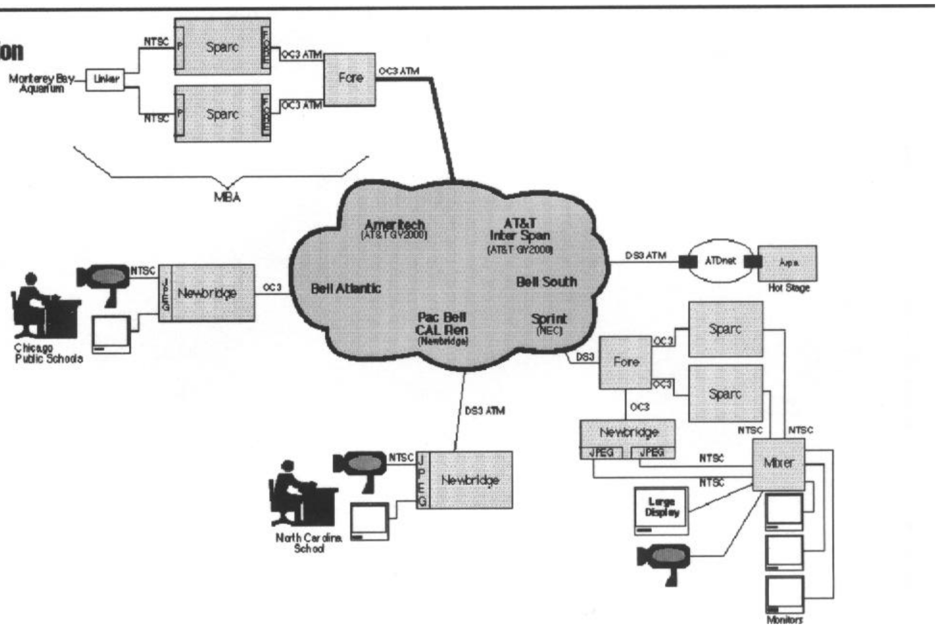

Chicago. Again, these Suns have parallax boards with Paradise software that make it possible to transmit video and audio across the network from one end to the other. Problems might occur because Argonne uses a higher-speed OC3 link while the University of Chicago uses a lower-sped DS-3 link. As a result, some traffic shaping must be performed at the ends or else some cells might be lost in between. Nevertheless, this particular network operated very nicely. Argonne's National Laboratory developed an accelerator known as the Advanced Photon Source (APS), and it is one of the newest accelerators in the world. The laboratory recently held its yearly users group, and while this very visually intensive information was being presented at the users' group, it was also transmitted back across the network to the University of Chicago. This is the location of several Of the APS's collaborative access teams which will be utilizing beam ports on the accelerator. That is one type of video application that has been put on MREN.

FIGURE 3

MREN Configuration

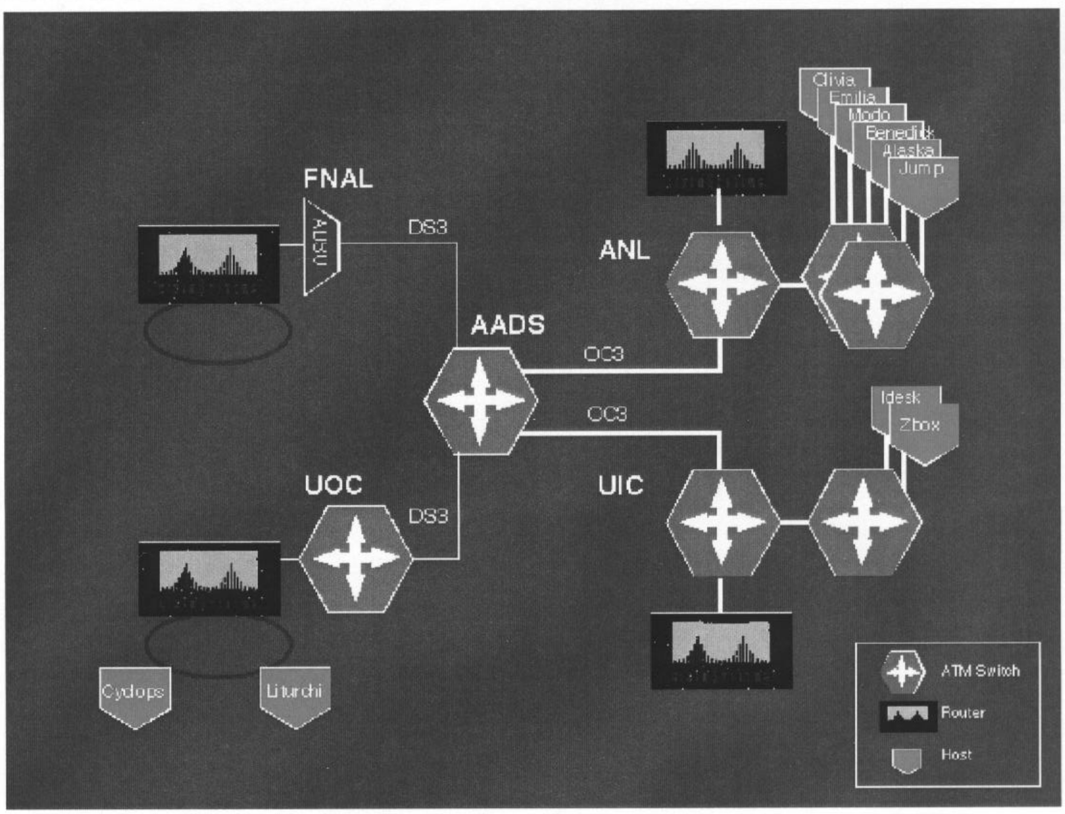

MREN Videoteleconferencing

There is also a need for Workgroup conferencing, the hadi-tional type of videoteleconferencing that accommodates people at various universities and research laboratories which want to communicate with people at Argonne in workgroup sessions. A multipoint control unit (MCU) is needed in this situation to allow multiple organizations to connect. ATM connectivity, an easy-to-use interface, and the highest-quality video are essential for the very visually intensive material that people discuss in these sessions. For this reason, MPEG-2-quality equipment is required. Many systems were evaluated, including the IBM WAVE equipment. The configuration that was tried out did not have a multipoint control unit, but an encoder/deencoder was placed at the University of Illinois at Chicago and another one was deployed at Argonne. This provided an OC-3 connection which was ATM-switched end to end, since both facilities had ATM switches. Both video and audio cameras were used. The architecture shown in *Figure* 5 was used both for an open how demonstration at Argonne as well as for some other scientific projects As in the previous figures, the ATM equipment is multivendor, which Argonne prefers to use.

Medica Imaging

Another application that Argonne National Laboratory is starting to get Involved with is medical imaging. The University of Chicago operates Argonne, and it has a fair-sized medical program at the university. One of their areas of Interest is medical imaging. Researchers at the National Institutes of Health (NIH) who are setting up another medical-imaging system are more interested in clinical applications, whereas the university is more interested in research applications.

At the NIH, there are a number of workstations spread across a wide-area network, or at least a metropolitan area network. They are transmitting very high-quality Images, usually images of X-rays. The network also supports videoteleconferencing so that a number of doctors can see the same high-quality Image. They can discuss, for example, whether they see a tumor or some problem in the X-ray. Whiteboard capabilities are also available so participants can circle and talk about various part5 of the image. This again features very high-quality images and very interactive videoteleconferencing.

Work in the area of radiation oncology at the NIH and at the National Cancer Institute uses this type of medical imaging videoconferencing. Some of the other applications are in the areas of cardiology, nuclear medicine, and monography. In the field of cardiology, instruments are now available that can actually take moving pities of the heart as it is beating. So, doctors can view actual video in addition to high-quality still pictures. In view of the important medical decisions that must be made in this environment, the video images must be of the highest quality. In the future, many experts will be able to consult at the same time, whereas now an X-ray typically must be shipped across the country for somebody else to look at.

Digital Video Servers

At Argonne, there is a research project called the System Voyager, whose purpose is to develop a next-generation hypermedia server architecture. It is a digital server based around a 12-node IBM SP2 parallel computer. The server puts out multiple streams, but the interesting aspect is not so much that stored visual images can be placed on a server for later transmission. What is interesting is the environment for making and presenting multimedia information that will be put onto the server. What are being developed are essentially collaborative environment tools that enable people in the areas of science and education to obtain information in a collaborative and useful manner. This means developing tools with which one can easily search or index material that is stored on the server. It means extending this to an environment much like the current World Wide Web, in which documents can be hyperlinked to media.

The idea is to build and link virtual organizations and get away from the present model in which people come to work in a physical laboratory There are some very expensive re-

FIGURE 4

MJPEG Over IP/ATM

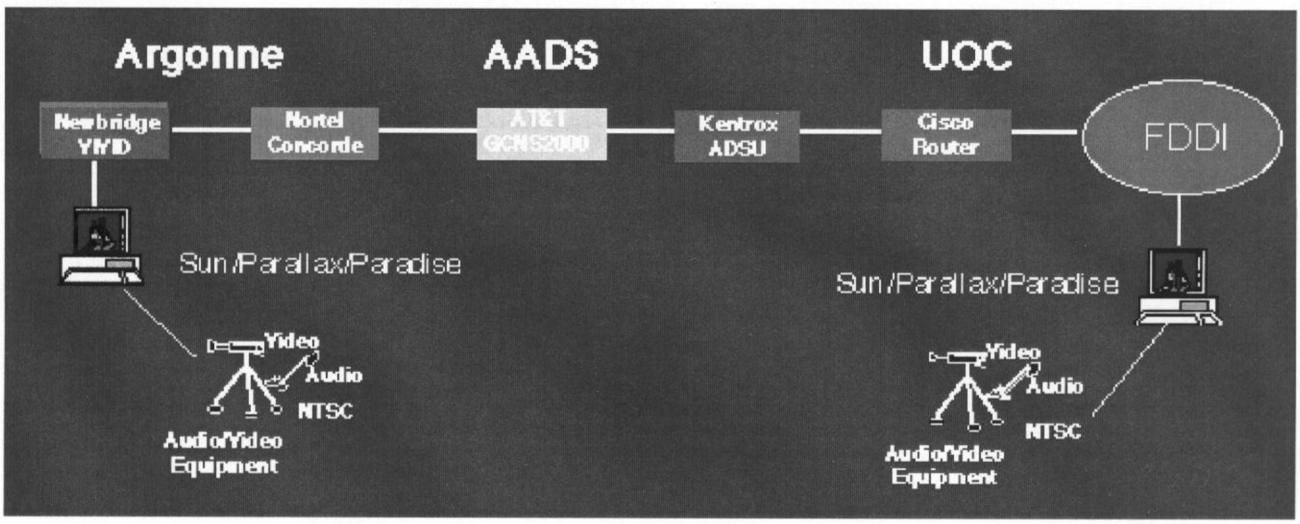

FIGURE 5
MPEG2 Over ATM Experiment

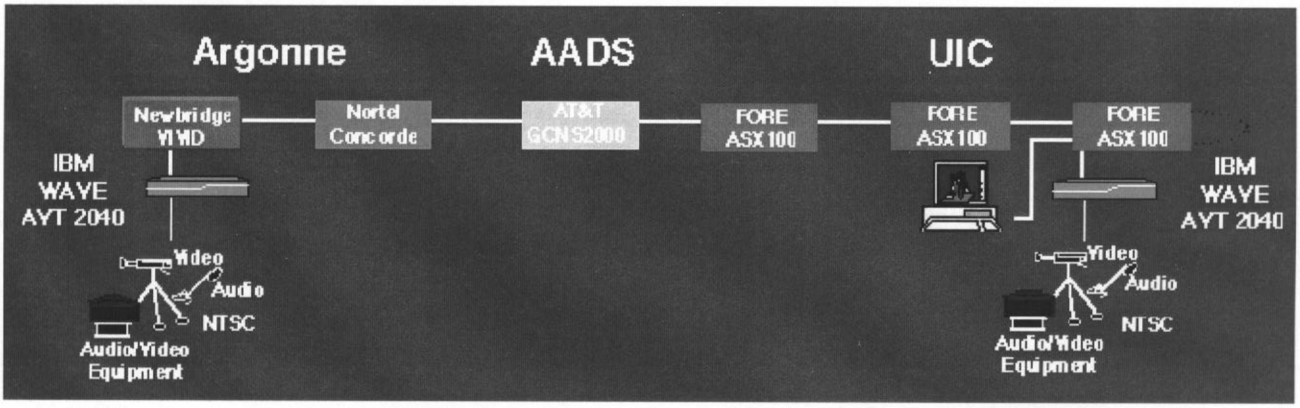

search instruments at the various laboratories and universities across the country, such as Argonne's Advanced Photon Source Accelerator (APS). The goal of Voyager is to enable researchers from around the country who are collaborating on the same project to link together across a high-speed network in this virtual environment. Such a setup would allow individuals to get information off the digital server and then view and discuss it at the same time. Argonne is building an environment in which it is easy to make and store digital information. But, more importantly, the laboratory is building collaborative tools around its digital server.

A Virtual *Reality Application*

The last sort of application I would like to discuss is the virtual reality CAVE, which was briefly described at the beginning of this paper. The picture that was taken of a simulated model of a finite element analysis grinding process is shown in Figure 6. It shows a bii wheel grinding on some sort of block. It is modeled and runs on Argonne's larger IBM SP2, a 128-node parallel processor array. The output of the model is then transmitted to the Silicon Graphics computers that drive the CAVE. A model is then produced that a person can view with glasses and actually become immersed within. It is possible for a person to stick one's head into the grinding wheel and see what the center of the wheel looks like. Of course, the important part here is that the model shows temperature distributions across the display. The model is capable of actually showing distributions within the grinding wheel, although they have not quite implemented that into the CAVE model yet.

The purpose of this system is to interact with the model. S-3, when inside the virtual reality CAVE, the user can move the wheel up and down, make it spin faster, put it deeper into the block, etc. A person can essentially study the wheel and see what happens under different circumstances. This is just one example of a virtual reality CAVE application.

The applications in this area are endless. The next step will be to link several virtual reality CAVEs together. There is one at

FIGURE 6
MPEG2 Over ATM Experiment

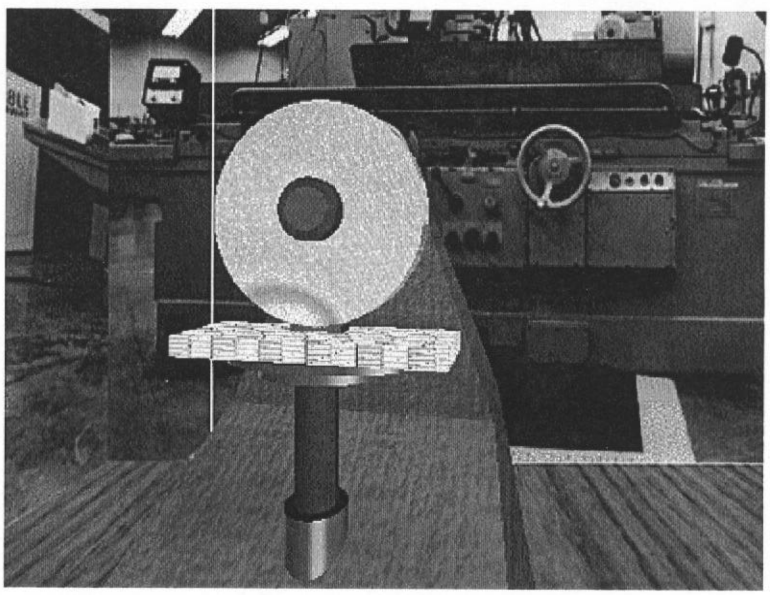

Argonne, one at the University of Illinois at Chicago, and one at the National Center for Supercomputer Applications (NCSA) in Urbana, and CAVEs are starting to pop up across the country. They can be linked together in order to build a virtual organization in which researchers see the same model, interact with it, and discuss it. In addition to this, there will be a videoteleconferencing component so users can see people at the other sites and talk to them. As with medical imaging, this involves transmitting very intensive images and being able to have videoteleconferencing as well.

Digital Video Technology Vision

The vision of the Argonne National Laboratory for digital video, particularly in the research and education environment, includes the following:

1. very high-quality digital video, like MPEG-2

2. a high-performance network infrastructure—essential for connecting remote sites together ("high performance" means high-speed ATM and SONET rates of hundreds of megabits per second)

3. the ability to interlink, interconnect, and visualize the various components of the various organizations—such as supercomputers

4. virtual reality systems—(There is something called an "immersadesk" which is a very large, desk-size version of a virtual reality room).

5. the ability to hook together mass storage systems. At Argonne now, there is a 35-terabyte mass storage capability, which represents a great deal of information. This is used to store information (such as astronomical observations) which involves huge amounts of data. In addition, NIH has many scans of the human body. This information is being stored so that a virtual reality model of a human being can be built.

6. support of a very collaborative environment, which includes the interconnection of digital video stores, mass storage systems, supercomputers, virtual reality CAVEs, immersadesks, etc. The system should provide for videoteleconferencing across a high-performance network, and it should allow for some very high-quality digital video.

   All these components play an essential part in building the virtual organization.

Electronic Commerce Market Demand Versus the Potential of Loop Access Technologies

Stephen L. Barreca
Vice President—Communications Technology Strategies
Technology Futures, Inc.

Serious problems arise when new technologies or markets threaten the core markets of a business, especially when emerging technologies tend to progress away from the core strength of the business. History tells us that many corporations do not survive significant shifts in their core markets away from their basic strengths. The electronic commerce (EC) access market represents a major shift in the core market of local exchange carriers (LECs), and that shift, to a large extent, is outside of LECs' traditional core strength: local exchange service. For this reason it is very important for the telecommunication industry to understand these markets and plan specifically for them.

The purpose of this paper is to establish a foundation for electronic commerce access providers (ECAPs) to build a strategy that addresses the emerging electronic commerce markets and the technologies that will sustain this market. It will examine the consumer EC markets from the perspective of ECAPs and then contrast market expectations with the long-term capabilities of competing loop access technologies.

The Electronic Commerce Market

To an ECAP, electronic commerce services are any services that provide data access to customers. Data access is any communication access between remote applications. All applications that require data communication access are EC applications—the author excludes broadcast TV but includes computer-based multimedia applications.

This paper concentrates on the residential EC market, which represents the lion's share of this market. The findings presented, however, are also applicable to the small business sector, although the trends must be slightly shifted in time.

Loop Technology Strategy for EC

Maximizing accessibility and capability is the right strategy for participation in the emerging EC market. The ECAP's strategy must deliver expected capabilities where needed and for an acceptable price. This paper will evaluate emerging loop access technologies according to their ability to achieve this strategy. (This paper does not address the issue of price or cost directly; it is inherent in the expectations for accessibility).

Bandwidth is the principal issue concerning the capability of loop access technologies. It is perhaps the most significant consumer expectation, but not the only one. Other technological issues, while they do not present bottlenecks at this time, could become significant. Perceived limitations/capabilities of various access technologies, regardless of whether they are real or not, may become key differentiating factors. A competitor's strategy, for example, may include promoting selected technological issues to the forefront of customer expectations. AT&T's competitors successfully utilized such a strategy in the early 1980s: they made "all-digital" PBXs a major selling point. AT&T, whose PBX family was not all-digital, lost considerable market share to its competitors, even though AT&T's PBXs were just as capable and often had many advantages over the competition.

EC Market Drivers

There are three key market drivers: personal computers (PCs), PC applications, and the Internet.

Personal Computers

The EC market has evolved as a direct result of the proliferation of PCs and PC applications. The size of the residential EC market is determined by PC adoption in the home. The pace of PC technology advancement and the evolution of PC applications' performance capabilities will govern future consumer expectations. To understand the future evolution of the EC market, we must first have a grasp of the future evolution of PC technology.

PC Performance

Projections for PC performance in terms of processor speeds are shown in *Figure 1*. Historical performance for both Intel and non-Intel microprocessor speeds are separately trended and contrasted against the historical trends of traditional

FIGURE 1
Computer Process or Performance

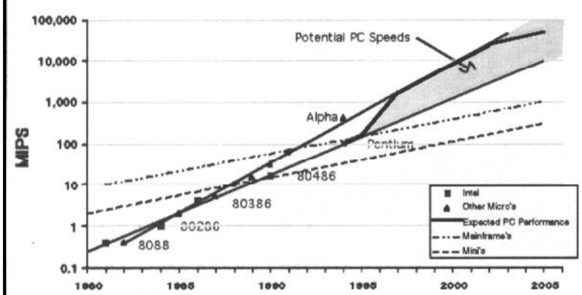

TABLE 1
PC Peripheral Performance

Technology	Years for 10x Improvement
Processor speeds	6
RAM size	6
I AN speeds	5
POTS modem speeds	a
Monitor resolution	10

mainframes and minicomputers. Essentially, the two microprocessor performance trends are very similar. While the historical trend for Intel may be somewhat less than that of others, increased competition in the PC processor market is predicted to cuase Intel to accelerate its performance. Signs of this acceleration are already occurring.

Looking ahead to 2005, PCs will run anywhere from 100 to 1000 times faster than today's PCs and will also cost considerably less. There are no technological barriers to achieving these speeds. Processing speeds may slow down if they exceed the demands of applications, but this is not likely. PC applications tend to gobble up all available speed, memory, and bandwidth.

In addition to providing insight into future microprocessor speeds, *Figure 1* also provides insight into the future performance of PCs in general. All related and supporting PC technologies are keeping pace with processor improvements. RAM memory, monitors, data storage devices, etc., are evolving at a pace sufficient to sustain the rapid advancement of PCs. Table 1 summarizes the historical pace of advancement of several PC-sustaining technologies.

At the same time that PC technology is rapidly advancing, the cost of PCs are rapidly declining. Ten to twelve years ago, a high-end PC and accompanying software cost around $10,000. Today, a comparable system costs about $2,500 to $3,500. IBM and others are pushing hard to roll out a $500 PC designed specifically for EC access, and PC-type adapters that allow Internet access using a TV are available today for less

FIGURE 2
PC Home Adoption

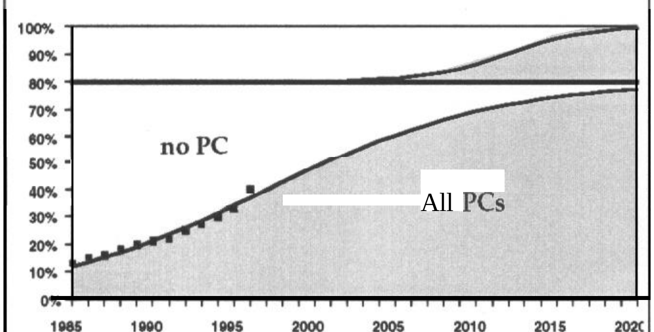

than $500. Shortly, PCs will be virtually free to the bulk of the home market-that is, cost will no longer be a significant factor in the purchase decision. (Five years ago, the author predicted that PCs would be virtually free by the year 2000-a prediction that was received with some amusement. Now it appears as if this estimate was overly conservative).

PC Adoption in the Home

Figure 2 represents a traditional view of PC penetration into US. households. PC home adoption is divided into two market segments. The first segment represents the 20 percent of U.S. households that are not likely to directly purchase a PC. The author anticipates that this market will acquire a PC when PCs are virtually free-i.e., integrated into the TV, telephone, or some other device, or when they are given away as part of subscribing to an EC service. PCs will begin to penetrate this 20 percent market near the year 2003, when PCs are virtually free, and saturate the market between the years 2015 and 2020.

The remainder of the market (80 percent of U.S. households) represents those households that are, or soon will be, willing to purchase PCs. The adoption projection shown in *Figure 2* is simply a best-fit substitution curve to the empirical data. For over ten years this curve reasonably fits the data; however, in recent years the adoption rate is slightly accelerating.

The recent increase in the adoption rate of PCs is likely due to the introduction of inexpensive multimedia PCs. They represent a competing technology to older and less robust PCs, and the markets have reacted. The combination of increasing power and functionality coupled with decreasing cost is leading to more rapid purchasing.

Figure 3 represents a departure from the traditional view of home PC adoption. Multimedia PCs are treated as a competing technology to older-type PCs. For the purpose of assessing the impact of multimedia PCs, U.S. households are divided into three segments or markets. Market 1 represents those households that owned a PC prior to 1993. Non-multimedia PC households peaked in 1993 at just under 28 percent of all U.S. households and has been rapidly declining since. This decline is due to the substitution or upgrading of older PCs with multimedia PCs. Continuation of this substitution suggests that multimedia PCs will replace virtually all older-type PCs by the year 2000.

Home PC Adoption—Multimedia

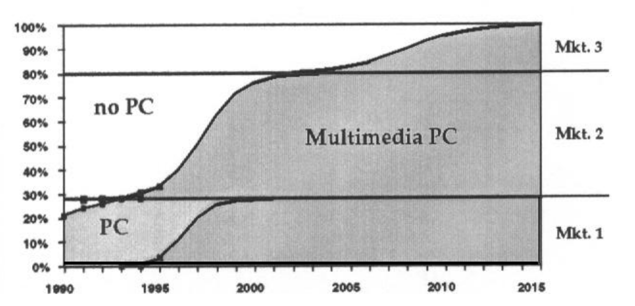

Market 2 in Figure 3 accounts for new PC purchases that are multimedia PCs. This segment represents 52 percent of all US. households. Continuation of the empirical substitution suggests that multimedia PCs will completely penetrate this market segment near the year 2003. The final segment, Market 3, represents the 20 percent of households that are not likely to directly purchase a PC, which was discussed in the previous section. PCs will start penetrating this market near the year 2003, and these PCs will be multimedia-type computers. It should be noted that data for only a few years of multimedia experience exists; nonetheless, these multimedia projections are reasonable and supported by the evidence.

The combined percentage of U.S. households with at least one PC is given in *Figure 4.* Since a PC or PC-like device is necessary for access to EC services, residential PC penetration provides an upper-bound on the size of the residential EC market.

PC Applications

Virtually every PC application would be improved with easy access to remote data and service sources; and virtually every organization, business, and service would benefit from direct access to its customer base and vice-versa. This is the reason behind the explosion of the Internet today. Furthermore, many household appliances could be improved with remote access capabilities, such as remote monitoring, maintenance, and control. The list of potential EC applications is endless and beyond the scope of this paper.

At the same time that PCs become virtually free, so will many EC applications. For example, AT&T is already giving away Internet access to its long-distance customers. Although most banks charge a fee (sometimes indirectly) for electronic banking services, this may soon be free; in fact, banks may even begin to pay the customer for using electronic access. Some providers of EC services have expressed a willingness to offset a portion of the costs to provide EC access.

FIGURE 4

Home PC Adoption

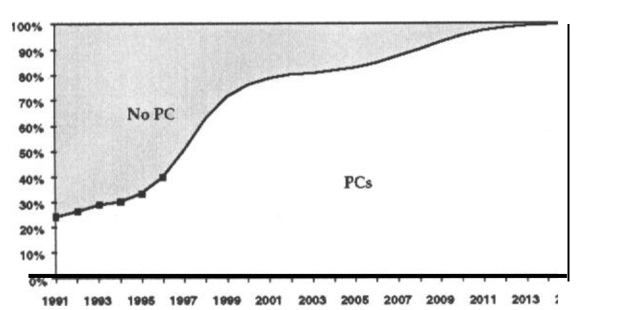

The Internet

Today, the vast majority of EC services are accessed via the Internet or via the Internet service provider's (ISP) network. Internet and ISP growth has been explosive, increasing the demand for Internet access. Rumors have circulated about the Internet having a meltdown. Although this is possible, it is unlikely to happen-this is not to say that tomorrow's Internet will look like today's. The Internet is not the only access service, and it will not be the only access service in the future. For example, financial institutions and other EC service providers will provide access to their servers directly via the public switched telephone network (PSTN) and perhaps other networks. Consumers will be able to dial into these servers and bypass the Internet (unless one considers other networks a part of the Internet).

The PSTN is an integral part of the Internet. It is the first network that residential and small-business customers encounter when going on-line. As sad as it may be, dial tone is the PSTN's "home page." However, what ISP would not want to be in that position? In the context of this paper, the PSTN should be considered as a vehicle for providing access to EC services for households and small businesses.

EC Market Size and *Growth*

Of major concern to ECAPs is the size and growth of the residential consumer market seeking access to EC services. The total EC market consists of U.S. households that access any type of remote service. U.S. households that use modems-a device necessary for full access is an accurate measure of the EC market today.

The first step in quantifying the residential EC market is to establish the size and growth of households which have PCs; this was done in the previous section (see Figures 3 and 4). Next, the percentage of these households that own modems must be determined. Finally, the size and growth of those households within this segment that actually use their modems needs to be assessed. This final segment quantifies the size and growth of the residential EC market. Companies interested in loop access technologies will want to weigh their capabilities against the demands of this market over time.

Figure 5 builds on the percentage of U.S. households with PCs for the three market segments reflected in Figure 3, to project PC households that have modems. Again, the consumers in market 3 will receive a PC-type device when it is given to them. The assumption is that this device will be an integral part of some service to which the consumer subscribes; thus, it will be equipped with communication capabilities. All PCs in market 2 (new PC households), by definition, will be multimedia devices; as such, they will be equipped with communication capabilities.

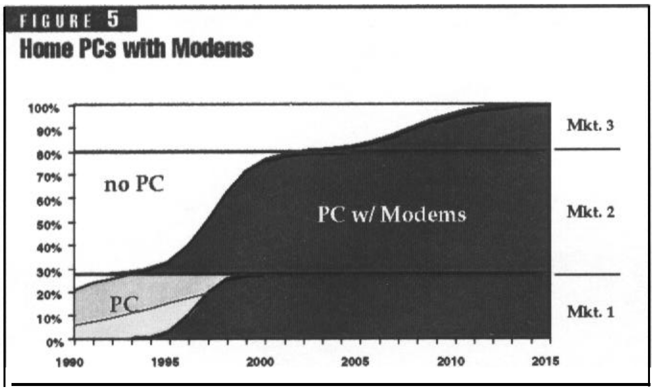

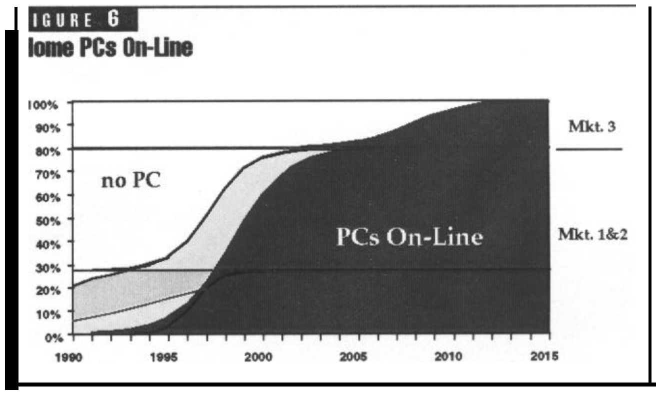

A significant percentage of the embedded base of PCs (market 1) do not currently have a modem. Continuing the empirical trend of modem penetration into this household market suggests that nearly all PCs will be equipped with modem capabilities near the year 2000. Modem penetration, if and when it is accomplished, does not Imply modem usage; today a significant number of households with modem-equipped PCs do not use the modem. To qualify as an EC consumer necessitates usage of the modem. Therefore, what is of interest are the number of households who both have and use modems.

Figure 6 presents the percentages of home PCs that are on-line (i.e., use modems). Given the nature of market segment 3, it is assumed that all household PCs in this market "se their modem. The universe of potential modem users in market segments 1 and 2 is the total number of households that have modem-equipped PCs. A substitution analysis, pitting households with modems versus households using modems, was performed on market segments 1 & 2 combined. The resulting trend shown In *Figure 6* suggests that-households using modems will increase from less than approximately 10 percent today to nearly 100 percent as early as 2005.[1]

That is, by 2005, all households with modem capability will be using it. This represents just over 80 percent of all U.S. households. This Is the consumer EC market. It is an extremely large and rapidly growing market. This is the market that ECAPs and ECAF aspirants are positioning to capture.

The explosive growth of the on-line market will significantly change consumer expectations of the PSTN and loop access technologies, and the implications of this are significant. Consumers will expect, for example, higher reliability, more bandwidth, and faster call completion, andcall duration will significantly increase. However, the requirement for bandwidth is the focus of the present paper.

Bandwidth Evolution

What will the future consumer EC market demand from the telecommunications industry in terms of bandwidth? Four to five years ago, Technology Futures, Inc., studied a host of EC services and came up with a view of the EC world. At that time, TFI subdivided EC applications into three categories: narrowband, wideband, and broadband services. TFI forecasted a gradual transition from narrowband to wideband and finally to the broadband environment. This was a very valuable projection because it projected the emergence of the EC market and quantified the gradual transition to higher-speed services.

Today, the definitions between narrowband, wideband, and broadband services are blurry. Just as the EC markets themselves are gradually evolving over time, so too are the bandwidth requirements. For many reasons, including the need for social and cultural changes, consumer expectations will not make a quantum leap from 28.8 kbps today to 45 Mbps tomorrow. Instead, bandwidth requirements will evolve commensurate with access technologies and capabilities of the public networks.

To equate EC market evolution to Increasing bandwidth requirements, the author subdivides the residential EC market into three segments:

1. Mass Market. This represents today's average consumer, who currently expects a 28.8 kbps modem to work over the telephone line.

2. High-end Market. This is today's high-end consumer who is willing to pay extra for 64-128 kbps communications.

3. Broadband Market. This is today's residential and small-business customer who is willing to pay a premium for 1.5-Mbps communications.

EC Mass Market

The EC mass market represents the average on-line consumer. Today, this market represents perhaps as much as 95 percent of the total EC consumer market. The appeal and benefits of emerging EC services increases the willingness-to-pay of many consumers, pushing them out of the mass market and into the high-end market. The relative size of the mass market will, therefore, shrink; it will settle at about 50 percent-70 percent of the total market.

Figure 7 projects the bandwidth requirements of the EC mass market. This data reflects the point in time when a particular modem rate becomes generally accepted by the average on-line consumer. Currently, these consumers use a 28.8 or 14.4 kbps modem. Those in the mass market use their existing telephone line to access some on-line service; they may be paying for the on-line service, but they do not pay extra for

FIGURE 7
EC Mass Market – Bandwidth Requirement

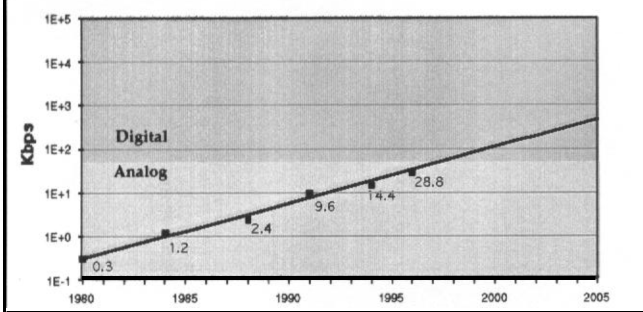

FIGURE 8
EC High-end Market–Bandwidth Requirement

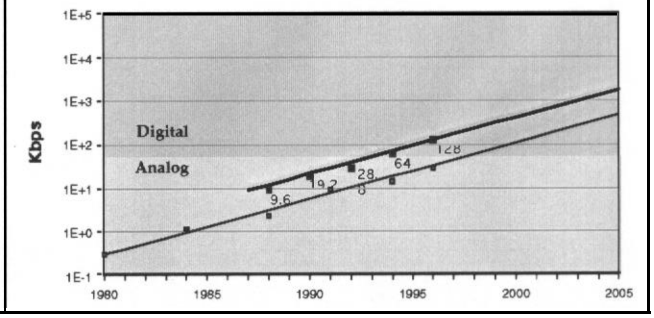

their access. The performance trend suggests that consumer demand for access speeds in this market segment will increase by an order of magnitude every eight years.

EC High-End Market

The EC high-end market is defined as today's consumers who are purchasing 56–128 kbps access. Presently this market represents less than 10 percent of the EC market. Over time however, the appeal of emerging on-line services, such as the Internet and other EC services, will encourage mass-market users to migrate to these higher bandwidths. The relative size of the high-end market will likely settle around 30 percent to 60 percent of the total EC market, but the percentage could be much greater if EC services yield enough benefits to warrant the added cost.

Figure *8* shows the projected bandwidth requirements of the high-end market. This market consists of customers who are willing to pay a little bit extra for access, such as an integrated services digital network (ISDN) connection. This market is paying $30 to $75 for that kind of access, and $15 to $25 for their on-line service. Individuals in this market have already made the transition to digital modems.

Four years ago, the high-end market was characterized by bandwidth usage in the range of 19.2–64 kbps. Four years before that, the bandwidth was 9.6 kbps. The high-end market has led the mass market by about four years, in terms of bandwidth. Continuation of this trend establishes a reasonable projection for the evolution of consumer bandwidth expectations in the high-end market; however, new technological developments that offer the potential of much higher modem speeds in the same price range could significantly increase consumer expectations.

EC Broadband Market

The EC broadband market is defined as today's 1.5 Mbps consumer (see Figure 9). Higher-speed users (DS-3 and greater) fall outside the residential and small-business consumer EC markets and are not addressed by this paper. The broadband market represents less than one percent of the total market today. Over time, a slight increase in the relative size of this market is possible, maximizing at about 5 percent-10 percent of the total market.

In both the EC mass market and the high-end market the bandwidth performance trends are very similar. It is reasonable to assume that a similar trend is likely for the broadband market as well. Again, as in the high-end market, new technological developments that offer the potential of much higher modem speeds in the same price range could significantly increase consumer expectations.

PSTN Bandwidth Evolution

The bandwidth evolution of the PSTN (shown in Figure 10) is the same as that of the EC mass market. The trend in Figure *10* is a simple, exponential progression, and there is reason to believe it will continue. The problem that occurs is that, to date, consumer bandwidth expectations in the mass market have been met with analog modem technology over the PSTN. Going forward, however, the current PSTN architecture and technology establish an upper limit on analog modem rates. Some suggest 33 kbps is the practical limit, while others believe 56 kbps is feasible. In either case, within two to four years, analog modem technology will have reached its practical upper limit over the PSTN.

If the PSTN is to be the principle vehicle of basic access service, a move from an analog to an all-digital n&work is necessary. The current EC mass-market consumer will expect future modem speeds as reflected in Figure 7 and *Figure 10,* if not higher speeds. Additionally, they will expect these levels of service for free, as part of their basic telephone service. To accommodate demand for basic telephone service, therefore, LECs must migrate to an end-to-end-digital network.

FIGURE 9
EC Broadband Market–Bandwidth Requirement

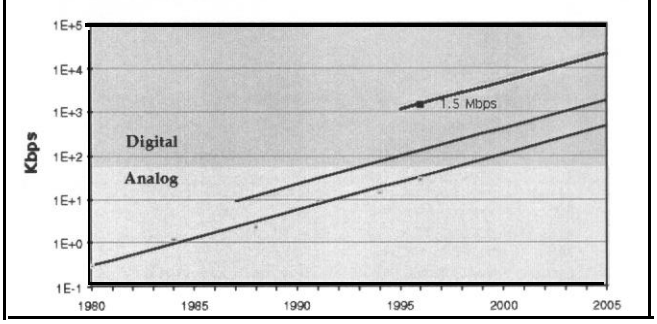

FIGURE 10
PSTN Bandwidth Evolution

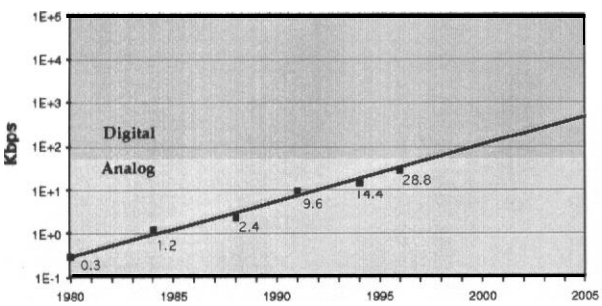

FIGURE 12
'SIN&P-P HDSL/ADSL

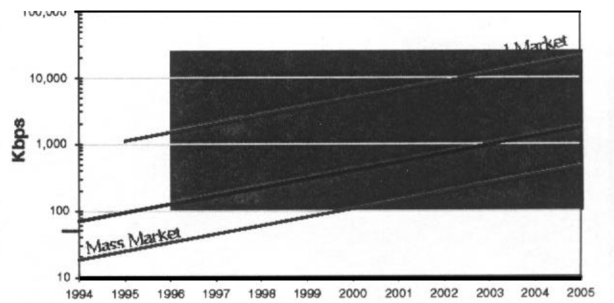

PSTN and ISDN

The shaded region in *Figure 11* shows the bandwidth capabilities of ISDN over the PSTN and is graduated to reflect the viability of ISDN to meet the demands of the EC market. The darker the shading, the greater the potential of ISDN to meet consumer expectations.

ISDN currently can deliver a basic rate of 128 kbps and a primary rate of 1.5 Mbpr Basic rate ISDN presently offers bandwidth suitable for the mass market; however the cost is above the mass market's willingness to pay. Nevertheless, ISDN costs are dropping significantly for the first time, and it is possible, although not likely, that within a few years these costs could be low enough to attract some of the mass market.

While basic rate ISDN is well positioned for the high-end market today, it is a fixed-bandwidth service. Basic rate ISDN (128 kbps) does not have the bandwidth that this market will demand in the future. The bandwidth trend of the high-end EC market reaches the bandwidth of primary rate ISDN (1.5 Mbps) near the year 2004. Additionally, it is reasonable to expect that cost decreases will be adequate in this time frame. Thus, primary rate ISDN could be suitable for the high-end market by 2004; however, in this time frame other emerging technologies will likely prove more advantageous.

While ISDN has a limited window of suitability, it should not be totally discounted. It is a very important strategic technology for any local exchange company. It is an international standard and is available throughout the United States. LECs are using ISDN to meet current and near-term demand while other technologies are being pursued.

FIGURE 11
PSTN & ISDN

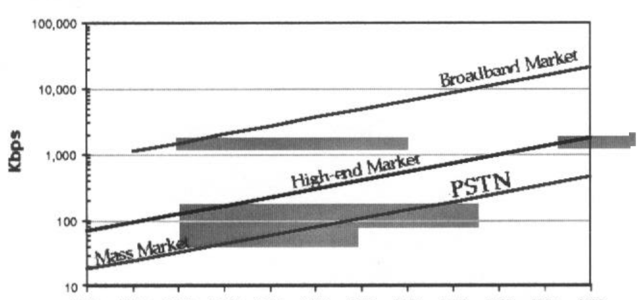

PSTN and HDSWADSL Technology Considerations

HDSL, ADSL, VDSL, etc.—collectively referred to as xDSL—are based on digital signal processing (DSP) technology, which is sophisticated software processing using microcomputers on a chip. There are 6–9 Mbps xDSL systems available that will work over the vast majority of a local exchange carrier's copper network. Their bandwidth capabilities exceed that of both the high-end and broadband EC markets (see Figure 12), and current cost indications suggest that xDSL EC access technologies cost less than ISDN.

The bandwidth requirement of both the broadband and high-end EC markets are currently less than the capability of xDSL technology. Additionally, xDSL is significantly less expensive than primary rate ISDN. Going forward, xDSL bandwidth capabilities will keep pace with EC market demand. In fact, the bandwidth expected by these markets will likely be governed by the pace of xDSL technological advancement.

To put the cost in perspective, a top-of-the-line ADSL chip set from Intel contains five processors on the chip, each of which runs at 50 Mbps (or 250 Mbps net). The chip set costs 5450 per 1,000. Chip sets from three years ago cost about $20 today. This suggests that today's top-of-the-line chip-set could sell as low as $5 in bulk, making today's device very attractive for the mass market in a few years.

The problem with ADSL technology from an access provider's perspective is that today it is basically a point-to-point (P-P) architecture. This is suitable for small niche markets, but it is too cumbersome for large-scale markets: trying to accommodate twenty million access lines with a point-to-point architecture is not practical from an administrative or operational standpoint. Hence, a more integrated xDSL architecture is needed.

If it evolves, xDSL-type technology will accommodate the broadband, high-end, and mass-market EC segments on a long-term basis. When xDSL becomes more integrated, service providers will be able to offer a very high-speed integrated digital network bundled with basic telephone service. Essentially, xDSL will be given away to the mass market. In short, xDSL has the long-term potential to satisfy the evolving electronic commerce markets.

FIGURE 13

PSTN: Impact of FTTC/FTTH

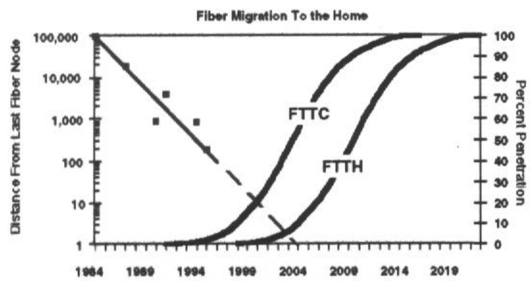

Source: Fiber Distance: Corning, Inc. 1994;% Fiber: BellSouth

Impact *of Fiber-in-the-Loop*

Another important question is what does fiber-t&he-curve (FTTC) or fiber-to-the-home (FTTH) do for the telecommunications industry? *Figure 13* shows how close fiber is getting to customers over time. The basic message is that by the year 2004, 10% of the PSTN will be served by a FTTH infrastructure.

FTTC is currently being deployed because it is more economical than copper for basic telephone service (both voice and narrowband data services). It is not being deployed to meet the demands of the EC market. Because xDSL's bandwidth capability is inversely proportional to the length of the copper loop, as fiber gets closer to the home and the length of copper in the last segment is reduced the bandwidth capability of xDSL Increases and the cost decreases.

The FTTC line in Figure 13 shows BellSouth's expectations for FTTC deployment. This projection is well on its way to being achieved. The author expects FTTH initial deployment to begin to vary slowly in 1998 and displace FTTC as the technology of choice a few years later. The author's projections and that of Corning Glass suggest that FTTH will reach about 10% penetration by 2004.

What will this deployment of fiber mean in terms of providing access to the EC market? Simply put, fiber allows companies to deliver more for less. Where FTTC is in place today, xDSL/fiber technology can deliver ten times the bandwidth at one-tenth the cost than xDSL/copper. FTTC allows a greater response to market demands, and it accomplishes this at a price much lower than copper. As a result, the emerging EC market and xDSL technology will accelerate the deployment of FTTC, forcing out copper faster than is indicated in *Figure 13.*

CATV Networks and Competing Technologies

Up to now, this paper has discussed telco strategies and access technologies. However, CATV companies' hybrid fiber/coax (HFC) network and cable modem technologies offer an effective strategy for CATV companies to get into the EC markets.

Cable-Modem Technology *and Usage*

Cable modems provide LAN-type access by establishing one or more LANs over the coax portion of the HFC network. An HFC/cable-modem system that serves 500 homes per node will support 10 Mbps transmission upstream and 30 Mbps using about 6 MHz of spectrum for each direction Ingress noise, originating principally from the subscriber, limits the application of upstream transmission to the bottom of the available spectrum (i.e., 5–54 MHz).

Use of a coax bus for LANs is an old technology. Today's high-speed office LANs are rapidly migrating to a point-to-point architecture over either copper twisted pairs or fiber. An HFC bus-star architecture will not support these newer high-speed LAN architectures. Continued advances in technology will slowly increase the bandwidth capability of HFC/cable-modem systems over time, and node sizes (currently around 2000 homes per node) can be reduced, which increases the available bandwidth per subscriber. Indications suggest that these advances will not keep pace with increases in demand over the long haul.

Availability and Cost

At the beginning of 1996, only LANCity had a commercial product on the market, but several vendors have since rolled out products. The typical customer premises equipment is a PC-LAN adapter. For the service provider, the main cost is in routers and servers located at the node. Overall cable-modem costs appear similar to that of ISDN. Although cost can be expected to decline, ever-increasing demand for greater bandwidth will likely offset these cost declines.

Bandwidth Limitations

Ingress noise from the subscriber, as noted above, limits the available spectrum to about 40 MHz. The theoretical upper limit for this spectrum is about 400 Mbps. From a practical standpoint, only about 40 Mbps is achievable with today's technology. Significant and costly technological advancements are necessary to yield appreciable increases in bandwidth

Currently the 10-megabit systems available are working in 6 MHz of spectrum. Overhead consumes more than half of this bandwidth, leaving about 4–5 megabits of useable bandwidth *Table* 2 illustrates the long-term potential of cable-modem systems in a symmetrical configuration (i.e., equal bandwidth up stream and downstream). By the year 2010, five 80-Mbps cable-modem systems will be required for each node, and node sizes must be reduced to 125 households just to capture half of the high-end market. Additionally, each cable-modem system will support only five simultaneous users.

While the example of *Table* 2 illustrates that cable-modems are bandwidth limited, it does not represent a definitive quantification of cable-modem potential. For example, an asymmetrical configuration that uses the 5–54 MHz spectrum for upstream transmission and higher spectrum for downstream is feasible and likely. On the other hand, *Table 2* is very generous in its anticipation of technological gains. Most experts do not expect that 80 Mbps upstream using only 6 MHz of spectrum is likely by the year 2010.

FIGURE 14
Cable-Modem Bandwidth Evolution

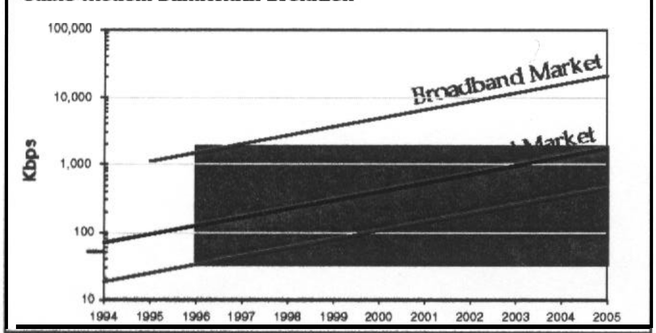

HFC Cable Modem Potential

Figure 14 projects cable-modem bandwidth evolution and how each of the three EC markets could utilize it. From 1997 to 2000, the state of the technology, the available spectrum, and the small market size will allow cable modems sufficient capacity to capture 50 percent of the high-end market with sufficient excess capacity to accommodate 30 percent-50 percent of the mass market (although price is a major factor for the mass market). Over this time period, the effective bandwidth per subscriber is less than half that needed by the broadband market; thus cable-modem penetration in the broadband market will be minimal.

Two main strategies must be pursued by HFC service providers to continue capturing 50 percent of the high-end market. First, cable-modem systems must be upgraded to state-of-the-art systems every four or five years. Second, the number of simultaneous subscribers per system must be continuously reduced. In spite of these steps, by around 2005 an HFC/cable-modem architecture will only be able to capture half the high-end market, leaving its excess capacity to go after the mass market. Beyond 2005, the technological maximum Mbps/MHz of cable modems is insufficient to yield the bandwidth required for the high-end market.

In summary, HFC/cable modems are an effective near- to mid-term strategy for HFC ECAPs. It is a good long-term strategy only if used in the interim while a FTTC/FTTH xDSL architecture is being built.

Competing Technologies: Wireless

Figure 15 presents one consultant's projections for the data bandwidth potential of PCS. If this estimate is correct, wireline carriers (LECs, CATV companies, etc.) should quit now, because PCS has won the game. In reality, wireless bandwidth today is insufficient for the EC markets. Ongoing technological advances in wireless bandwidth capacity simply will not keep pace with EC demand. There will be a market for mobile EC services, but it will be bandwidth-limited and therefore much more likely to be complimentary to the EC markets.

Summary

In summary, the EC market is evolving faster than previously documented. Multimedia PCs have accelerated the penetration of modem-equipped PCs into the home. The Internet and other on-line capabilities (work-at-home, financial transactions, etc.) are accelerating modem usage, and inexpensive access is making on-line usage affordable to the masses. The trends suggest that the EC market is on the verge of an explosion.

Potential Of Access *Technologies*

For the broadband EC market, primary rate ISDN and cable modems meet today's bandwidth requirement, but both are bandwidth-limited and offer insufficient bandwidth for the future. Only xDSL technology as an adjunct to the PSTN is feasible over the long term. Because xDSL technology exceeds the current and projected bandwidth demands of the broadband market segment-and at a cost lower than what the high-end market is willing to pay-the author expects these market segments to merge.

TABLE 2
Utilization of HFC/Cable–Modems for EC Markets

Year	Systems per Node	BW/Sys. (Mbps)	Simultaneous Subs/Sys	Effective BW/Sub	HH/Node	Market Share Obtained High-End Market	Market Share Obtained Mass Market
1996	6	10	10	0.450	500	50%	50%
2000	6	20	20	0.450	500	50%	34%
2005	6	40	11	1.640	500	50%	8%
2010	5	80	5	7.200	125	50%	13%

Market Objective

- Gel 50% of high-end EC market
- Ignore BB EC market (bandwidth too high)
- Use available spare capacity for mass EC market

Design Priority

- Bandwidth per subscriber must exceed 80% of market need
  1. Increase systems size (to the extent available)
  2. Decrease simultaneous users
- Get 50% of high-end market
  1. Split mde size as required

FIGURE 15
Cellular/PCS

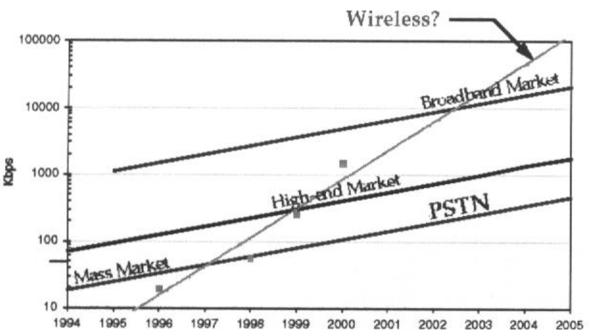

For the high-end market, cable modems give carriers with HFC networks a near-term strategy to play in the EC market while other alternatives can be developed. Basic rate ISDN gives LECs a short-term strategy to accommodate the high-end market while xDSL technologies are implemented. Point-to-point xDSL systems give LECs a medium-term vehicle with which to play in the market while more integrated xDSL architectures are developed.

For the mass market, the spare capacity of cable modems is available only in the very near term, but willingness to pay will limit cable modems in this market. Analog modems and ISDN give LECs a way to be active in this market for the short term, assuming ISDN prices reach parity with today's plain old telephone service (POTS) by 1998. For the mass market, xDSL is a long-term solution, if integration issues are resolved.

In short, there is no definitive long-term solution for LECs, CATV companies, or alternative LECs. An integrated, ubiquitous solution must be developed now for implementation shortly after the turn of the century. Today's potential ECAPs have technologies available to them for the near-term, but both must rapidly migrate to a longer-term solution.

Over time, market size will continue to grow rapidly, as will bandwidth demand. In this environment, xDSL has the greatest long-term market potential. Over the long haul, the overlay and point-to-point architecture of xDSL must migrate into a single integrated intelligent network.

The Killer Technology

At the beginning of this paper, the author stated that serious problems arise when new technologies threaten the core markets and strengths of a business. Such "disruptive technologies" are something that all Potential ECAPs must be on the watch for. If a new technology has greater advancement potential than the technologies sustaining a company's core business or its core strategy, then these companies better be mindful of the potential for disaster. Most companies do not survive this situation. The emerging loop access technologies are potential disruptive technologies.

The EC market will gobble up all the bandwidth that the consumer can afford. In other words, bandwidth demand is limited by affordable technology. This leaves a tremendous window of opportunity for a disruptive technology to emerge which blows the competition out of the water. The EC demand trends presented in this paper, while reasonable and supported by empirical data, are not immune to significant shifts resulting from disruptive technologies.

The emerging loop access technologies, namely xDSL and cable modems, are potentially disruptive technologies. Their unprecedented capabilities will undoubtedly shift consumer expectations. Pent-up demand and EC access devices (e.g., PCs) can absorb all of the capacity that emerging loop access technologies can deliver. The groundwork has been laid for a major shake-up in the markets and market providers.

The author's advice is hedge your bets; research all potential technologies; do not write off a new technological solution quickly; do not ignore the technological trends; and finally, do not believe that large incumbents have the advantage-IBM nearly went belly-up with the emergence of mini-computer technology. Events beyond IBM's control, such as the emergence of the PC market, helped to save the company.

Note

1. As mentioned earlier, the number of households on-line could be overstated in the long run because the substitution of multimedia PCs is only a few years old. At the same time, the number of PC households with more than one PC is growing rapidly as a percentage of all households with PCs, so the actual number of EC consumers is greater than the number of households using modems. This is not a concern in the near term, because the number of households that need simultaneous EC access is very small. However, ten years from now the author expects simultaneous access to be in very large demand.

From Internet to Superhighway: The Future of Interactive Services

Paul Budde
Managing Director
Paul Budde Communications

The Concept of Superhighways

Superhighways should be seen as a concept, rather than as a technology. This first became widely recognized during the United States' presidential elections in 1992. Al Gore used the concept to promote economic growth and referred to an earlier era when economic stimulus resulted from the building of the interstate (road) highway network. After the elections, Vice President Gore incorporated the concept in the National Information Infrastructure Initiative (NILL). The idea was to encourage private companies—notably the telephone and cable TV companies—to invest in this new infrastructure.

While there is a lot of talk about superhighways and even several trials and pilots, there are currently no commercial superhighways in operation anywhere in the world.

The superhighway concept is that of an integrated, broadband-network infrastructure delivering a wide range of digital, fully interactive, multimedia (i.e., voice, video, and all other data) services directly to homes, workplaces, institutions, and public places (See *Figure 1*).

The definitive factors are:

- *broadband*—widely available, reliable, high quality, and wide-ranging services catering to all possible tastes, providing lots of scope for individual choice
- *interactivity*—full, two-way interactivity means that services will be available on demand; individuals may tailor the services and contents to their own liking and provide feedback to service providers
- *digital*—in order to provide the capacity that the above elements demand, the network must be digital

The underlying technology to be built is fiber-optic cable, although recent developments in digital compression and in coaxial cable may provide enough bandwidth in the short term. Other technologies are also coming to the forefront. It can be seen then, that while several of the current—often hybrid—networks, some of which are mistakenly referred to as superhighways, meet a few of these criteria, none meet all. The Internet, for example, is an on-line, narrowband service; cable TV is based on forty-year-old analogue technology. However, both can be seen as ramps to the superhighways.

The Internet over telephone networks (28.8 kbps and 64 kbps) will dominate the market in 1998. From that period onwards, coax cable and fiber-optic cable networks are going to make an impact that will become more noticable around the year 2000. After that, growth of true superhighways will be based on fiber-optic backbone networks and asynchronous transfer mode (ATM) switching technologies.

Vision of the Future

The annual outlook for the telecommunications industry produced by the University of Southern California indicates that

FIGURE 1
Key Elements of Superhighways

- no mass deployment of superhighways until 2000-2005;
- media, not telecommunication, driven;
- not one infrastructure, but a range of roads and highways;
- hybrid technologies (I call them superhighways ramps) to test services, markets and applications will have to be built first;
- no mass media approach but niche market strategies to deliver individual products;
- hi-tech/hi-touch – clever usage of human beings);
- future offices will be more and more at home;
- more involvement from creative people, advertising, script writers, art directors, film producers.

FIGURE 2

Vision of the Future: Now–2025

Now	Now - 2000	2000-2015	2015-2025
Telephony.	Secure trading/networks.	Video-on-demand.	Fully interactive telecommunications/entertainment.
ISDN.	Multigagabit LAN/MAN.	Video browsing.	Fat panel 3D display.
E-mail.	High definition TV.	Terabit networks.	Real-time translation.
Interactive voice.	User controlled networks.	Multimedia networks.	Personal agents.
Low bit video.	Smart PCS terminals.	Speech recognition.	Fault tolerant networks.
Cellular mobile.	Deaf services.	Seamless language.	
	Electronic libraries.	Customised media.	

there will be a structural impact on the migration of the entire information industry towards digital electronics:

- the industry will reorient itself horizontally
- as major industry players position themselves for the future, there will be a series of within- and across-industry consolidations based on a chosen functional specialization
- the transformation will result in only three major industries, which suggests that it is going to become more efficient in the process; the three industries will be providers of digitized content, multimedia devices, and convergent networks

Therefore, companies that thrive in the future will have a bias toward personal rather than institutional markets; greater expertise and experience with digital electronics; a functional edge, i.e., outstanding expertise in performing one or more information functions; and more experience with video-based information—technologically, other forms of information are a subset of video (see *Figure 2*).

In addition to the functional basis, three key factors characterize the evolution of the information industry in coming years each new industry will be multimedia in nature, given the convergent nature of digital electronics; each of the sectors will be global; and the industry will become increasingly driven by a model in which personal markets rather than institutional markets will become the lead markets for technology deployment.

Of the three primary sectors of the new information industry, the content area has advanced furthest towards the new model. First, a great deal of content is already digitized and ready to be used in multiple ways. Second, the rationalization of content businesses has been underway for some time, so that we already have large content entities in the United States such as Time Warner (which may soon add TBS) and the Viacom-Paramount-Blockbuster combination.

International Superhighways Market Statistics

Technological trends are bringing the telecommunications and the broadcasting sectors closer together. In particular, on digital networks it is irrelevant to distinguish between transmission of images from voice, data or text. Consequently, public telephone operators are considering entering the market for two-way video transmission (e.g., video conferencing, video telephony, video on demand) as a way of exploiting the full capacity of the new generation of broadband fiber-based networks. For their part, cable TV companies are looking to provide higher value services—such as homeshopping, electronic information services, pay per view, and perhaps even telephony—as a way of gaining new revenue sources from their existing subscriber base. Inevitably, these trends are bringing the two sectors into competition with each other. Cable TV and direct-to-home (DTH) satellite services reach over 220 million households. Roughly one-third of these are installed in North America and Europe and 25 percent are in Asia (see *Table 1*).

Some of the world's biggest cable TV markets are in nonOECD countries. The second largest market in the world in 1995 was in China, which had around 35 million subscribers. India had 16 million subscribers, a market larger in size to that of Japan. Other large Asian markets include Korea, with more than 2.5 million and Taiwan, with 3.6 million. Two other countries are also worthy of mention—Argentina had 4.5 million subscribers in 1995 and Poland boasted 2.7 million cable subscribers.

Poland's satellite TV market is also quite significant. With around 1.8 million subscribers, it was the fifth largest satellite market in the world. Hungary also has a sizeable market—900,000 satellite and 1.3 million thousand cable TV subscribers in 1995. Still, in Eastern Europe, the Czech Republic had around 650,000 thousand satellite receivers. In Asia, sizable satellite markets are to be found in Indonesia and India both around 1 million in 1995.

Markets for Superhighways

Global Overview

The U.S. broadcasting market is currently as big as Europe and Japan together, although these two zones are displaying greater dynamism. Overall, the audiovisual sector is continuing its trend away from direct financing through advertising and license fees and towards direct financing by the consumer (see *Table 2*). This largely explains the priority the major groups are according to penetrating pay TV markets and the increasing importance of distribution in the audiovisual field.

The top world audiovisual companies (see *Table 3*) are showing a certain degree of stability, even if their cumulated turnover is rising rapidly (15.8 percent), the result of market

TABLE 1

Top Cable and Satellite TV Markets (Ranked by Size) in OECD Countries, 1995

Country	No. of cable subscribers (millions)		No. of satellite receivers (millions)
	1995	1994	1995
USA	62.96	60.5	3.8
Germany	15.81	14.6	9.5
Japan	12.30	10.4	5.9
Netherlands	5.70	5.8	0.3
UK	1.42	0.92	3.4
Belgium	3.63	3.6	0.25
Switzerland	2.40	2.3	0.2
France	1.50	1.6	1.1
Sweden	1.90	1.9	0.7
Mexico	1.14	2.0	
Austria	1.04	1.03	0.98
Finland	0.83	0.8	0.15

(Source: Paul Budde Communication based on data from ITU, OECD, Siemens 1997)

TABLE 2

Service Providers' Revenues Received from Interactive TV Subscribers and Services, 1995–2004 (US $ Million)

	1995	1998	2001	2004
Europe				
Basic subscriptions	3	373	2,600	7,310
Additional services	1	168	1,100	2,960
Advertising	1	181	1,210	3,350
Traffic	0	22	176	513
Total Europe	**6**	**744**	**5,100**	**14,100**
United States				
Basic subscriptions	6	290	2,590	8,960
Additional services	3	170	1,1440	4,690
Advertising	4	220	1,960	6,790
Traffic	0	21	188	647
Total USA	**14**	**701**	**6,170**	**21,100**
Asia/Pacific				
Basic subscriptions	0	30	274	1,020
Additional services	0	14	121	426
Advertising	0	16	152	565
Traffic	0	5	47	149
Total Asia/Pacific	**1**	**65**	**594**	**2,160**
All regions				
Basic subscriptions	9	693	5,470	17,300
Additional services	5	352	2,670	8,080
Advertising	6	417	3,320	10,700
Traffic	1	49	411	1,310
TOTAL	**20**	**1,510**	**11,900**	**37,400**

(Source: Paul Budde Communication based on data from Ovum)

TABLE 3

The Top 10 Audiovisual Companies Worldwide, 1995

Rank 1995	Company	Country	Total turnover 1995 - US$bil.	Audiovisual turnover 1995 - US$bil.	Audiovisual turnover 1994 - US$bil.	Variation AV to 1995/1995	Rank 1994
1	Time Warner	USA	17,696	13,770	12,495	15.3%	1
2	Viacom	USA	11,688	9,172	5,242	75.0%	7
3	Sony	Japan	47,619	8,618	7,815	7.0%	2
4	Tele-Communications	USA	6,851	6,851	4,396	38.8%	8
5	ARD	Germany	6,531	6531	5,650	0.2%	5
6	NHK (1)	Japan	6,043	6,043	5,744	2.1%	4
7	Walt Disney	USA	12,112	6,004	4,793	25.2%	9
8	Capital Cities/ABC	USA	6,679	5,728	5,277	8.5%	6
9	PolyGram	Netherlands	5,479	5,479	4,725	2.3%	10
10	Bertelsmann	Japan	15,029	5,127	n/a	8.4%	15

(Source: Paul Budde Communication based on data from IDATE France)

growth as well as concentra-tion of players. Viacom, after merging with Paramount and Blockbuster, has taken its place with Time Warner, Sony, MCA, NHK, ARD, and Capital/Cities among the companies with an audio-visual turn-over exceeding US$5 billion.

Other significant increases are those achieved by BSkyB, MAI, Granada, and Carlton, which testify to the dynamism and concentration present in the British audiovisual sector, and by the German Pro7 private TV company that boasted a 39 percent growth in turnover. Conversely, Nintendo, Sega and RTVE are among those that have recorded a net decrease in performance.

On the world scene, it is to be noted that European and American companies are of roughly equal market importance—even though they differ radically in nature. Germany dominates in Europe (accounting for 37 percent of the turnover of European companies among the top 100), followed by the United Kingdom with 22 percent, while France occupies a relatively modest place with 11 percent.

The major companies have admittedly recorded higher growth rates, but the same cannot be said for their profitability. The average rate of the top 100 companies is in fact only around 3.9 percent of turnover, compared with 4.3 percent the previous year. Those with the highest profitability include TBL (Hong Kong), BSkyB (UK), M6 (France), and News Corp (Australia).

The Ramps to the Superhighway

The ramps to the superhighways consist of current telephone, computer, and TV infrastructures.

While still in its early stages, the Internet is currently the most important development in the superhighway market. While the concept of the service has been around for over twenty years (on-line, videotex, teletext, EDI, etc.) it was not until the mid-90's that a breakthrough occurred in this market. The major reason for this was that the PC market had reached sufficient mass to move away from being a niche market to becoming a mass market.

In 1996, the home-computer market was larger than the business PC market (see *Table 4*). The global PC market is growing at a rate of around 15 percent per year.

Current Internet, cable TV, and other networks will, over time, be upgraded and replaced with superhighway networks such as full services networks (FSN). *Tables 5* and *6*, which provide statistics on the current infrastructures, give an indication of where things might be going.

Several countries have made commitments regarding their superhighways and they claim that these infrastructures are under construction (see *Figure 3*).

To even come close to delivering on these promises, a range of issues will need to be solved, including:

- current high costs of delivery, based on current or proposed network infrastructures;
- national and international interconnect fees, which are hampering competition;
- vested interests in markets that will have to converge are stalling progress; and
- lack of procompetitive regulatory environment nationally, World Trade Organization (WTO), Organization for

TABLE 4

Homes with Personal Computers

Country	Penetration/homes
China	7%
India	9%
Malaysia	20%
Japan	21%
France	22%
UK	25%
Germany	30%
Singapore	32%
Hong Kong	32%
Australia	35%
United States	39%

(Source: Paul Budde Communication based on data from Link Resources)

TABLE 5

Vehicles of the Global Information Highway

Television sets	59%
Telephone lines	32%
PCs	9%
Total global installation base	**2 billion units**

(Source: Paul Budde Communication based on data from ITU)

Economic Cooperation and Development (QECD), and General Agree-ment on Tariffs and Trade (GATT).

In short, the global superhighway will evolve along the lines of evolution, not revolution (see *Table 6*).

In 1994, we witnessed telecommunication carriers around the world frantically jumping on the superhighway bandwagon. Global media and telecommunication joint ventures were happening everywhere. The trend, fuelled by the mega-merger negotiations in the United States, spilled over into Europe and Australia where mergers were formed, many of which subsequently collapsed.

The telecommunication and communications industry generated 6 percent of the world's trading revenue in 1995.

Now, in 1997, the mega players still do not have a clear vision of where they are going. The same is true for government regulators and policy makers, investors, developers and just about everyone else involved. All of this demonstrates that driving on the superhighways is not easy and nerves of steel, lots of stamina and sharp vision are essential. There is no doubt that the rewards for the winners will be high, but there is little room for the faint hearted.

The key issues in this field are delivery technologies (availability of the new technologies, their often unknown commer-

TABLE 6

Top 20 Plus Countries Most Likely to Take a Lead in Superhighways

Multimedia access: main telephone lines, TV sets and PCs per 100 inhabitants / major economies in 1995			
Economy	**Phone density**	**TV density**	**PC density**
United States	62.6	78	32.8
Denmark	61.3	54	27.1
Canada	59.0	65	19.3
Sweden	68.1	48	19.3
Australia	*51.0*	*64*	*27.6*
France	55.8	58	13.4
Switzerland	61.3	46	34.8
Netherlands	52.5	50	20.1
Germany	49.4	55	16.5
Japan	48.7	62	15.3
UK	50.2	61	18.6
Austria	46.6	50	12.4
Belgium	45.8	46	13.8
Singapore	47.9	36	17.2
Hong Kong	53.0	36	11.6
Italy	43.4	44	8.4
Spain	38.5	49	8.2
Korea (Rep)	41.5	32	12.1
Taiwan	43.1	32	8.3
Hungary	18.5	44	3.9
Malaysia	16.6	23	4.0
Thailand	5.9	23	1.5
China	3.4	25	0.2
Philippines	2.1	13	1.1
Indonesia	1.7	15	0.4
India	1.3	6	0.1
Average high income	**53.2**	**61**	**20.5**
Average upper middle income	**14.5**	**26**	**3.3**
Average lower middle income	**9.1**	**20**	**1.1**
World average	**12.14**	**23**	**4.2**

(Source: Paul Budde Communication based on data from ITU)

FIGURE 3

Superhighways Under Construction

Australia	passing 7 million homes by 1997
Canada	all research and educational communities connected by 1999
Europe	open infrastructure by 1998
Japan	100% coverage by 2015
Singapore	fibre to the kerb by 1997
South Korea	fibre to all major cities and offices by 1997
United States	all schools, libraries, hospitals, etc connected by 2000

FIGURE 4

Superhighway Principles

- free access to information necessary for full participation in a democratic society;
- guaranteed freedom of speech for network users, protection of copyright for creators and the right of library patrons to use material;
- a vital civic sector component analogous to an "electronic commons" for public discussion and debate;
- a healthy marketplace of ideas that is accessible to all and not controlled by telecommunications carriers;
- a set of policies that protect the privacy of the users;
- a chance for the public to be fully involved in making policies related to the network.

cial potential and life-cycles); services and content (what will people want, need and buy?); and regulatory and social policy issues (see *Figure 4*).

Carol Weinbaus, Research Director of the University of Florida, United States, highlighted the following issues that the public policy debate needs to move on:-

• What are the rules by which multiple providers will be allowed to build broadband networks to serve customers?

• What are the public policy objectives for information superhighways? Will the networks be driven by customers or by new government policies for universal service? Will this involve subsidies? If so, what will they look like?

• How will the multiple networks interconnect to provide customers with services that are easy and inexpensive to use? What will be the interconnection standards?

• What should be the structure for intercompany payments to facilitate advanced networks and services?

• What is the appropriate regulatory model for multivendor digital networks?

Infrastructure Developments

Still a Long Way to Go

Superhighways are still more of a concept than a definition. Introduced in the United States as a National Infrastructure Plan, information highways can be generally described as high-speed communication infrastructures capable of delivering all types of information anywhere by means of friendly and possibly cheap interface, as shown in *Figure 5*. The U.S. Government's initiative was followed by Singapore, Australia, Japan, and France. By 1997, many other countries had launched similar plans. Most countries are—in one way or another—addressing information highway issues. If not on a superhighway level, at least Internet highways are being explored (see *Figure 6*).

Internet and ISDN: The Runner-Ups

A key driver behind the current global interest in superhighways is the success of the Internet. Especially as electronic trading on the Internet is secured and facilitated by appropriate financial and administrative service. The effects of the Internet will be felt far beyond e-mail and on-line services, it has the potential to change the total telecommunications infrastructure, whereby the carriers and operators are no longer the central players, but where the customers are in control. The Internet has the possibility to overtake the telephone and cable TV market in size within the next five years. For this rea-son, the more advanced telecommunication broadband networks are put on the backburner, in favor of high-speed data networks that can be installed more quickly. Too little, too late, ISDN is trying to take a leading role in this, finally, after more than twenty years, finding a market for its technology.

While the developments in technology might be mind boggling, the greatest effects of the new developments will be in intellectual property rights, remote electronic trading and financial services (e-cash), taxation issues, etc. Very few countries, companies, and societies are ready for the dramatic changes that are about to occur.

FIGURE 5

Key Elements of the Superhighway Infrastructure

- Speed offered by increasingly available bandwidth, new transmission techniques and the use of optical fibres
- Intelligent networks that are simpler and easier to manage and that can be used flexibly, dynamically and in a customised fashion
- Ubiquity and mobility of end user systems that is likely to come from miniaturisation and wireless communication

FIGURE 6

Key Superhighway Techniques

ISDN	Integrated Services Digital Network
	integration of voice and data services
B-ISDN	Broadband Integrated Services Digital Network
	standardised wide range of sophisticated broadband services
SDH	Synchronous Digital Hierarchy
	Transmission technology to 'police' digital flow of information
ATM	Asynchronous Transfer Mode
	The switching technology to allocate speeds and channel mix
Frame relay	Switching technology to improve current data services
MANs	Metropolitan Area Networks
	Networks that can be used in some ten square kilometres
FDDI	Fibre Distributed Data Interface
	High speed data transmission in LANs (Local Area Networks)
ADSL	Asymmetrical Digital Subscriber Line
	Interactive TV on traditional networks

(NB. These technologies are all described in more detail in the Information Technology Management Report and information modules)

(Source: Information Technology Management Report 1997, Paul Budde Communication)

Apart from the increase in speed needed to deliver the new services, intelligence will need to be added to the network. A range of new services will have to be delivered to different people in different formats. The early driver of intelligent networks (INs) has been freephone, followed by virtual private networks (VPN) and calling cards. Other IN applications include premium rate services, televoting, and personal numbering (see *Figure 7*).

Commitment Needed from the Carriers

Currently, the biggest problem with the Internet is lack of bandwidth, leading to extremely low transmission speeds. Even with the top of the current range modems and high-quality Internet servers at the systems level, the speed is too slow. Many first-time users are so put off by this that using the In-ternet becomes a one-time experience.

The only thing that can save the Internet is higher speeds. At the moment there is no way that telephone systems will be able to offer this. Solutions available within an acceptable period of time are cable modems and satellite TV modems; however, a new generation of ADSL modems might offer a solution for the telephone network operators. What is needed is at least a ten-fold increase in speed.

However, the technology is not so much the problem, nor is consumer acceptance. The question is, are carriers and service providers prepared to commit themselves to the technology? The lesson learned from ISDN is that a half-hearted approach is not going to work. If carriers and service providers are serious about ADSL, they will have to put the resources into it. Realistic planning and pricing are essential ingredients in such a plan as well. Telecommunication companies have failed to penetrate the market with their alternative, ISDN services. While this technology has been around for twenty years, they have kept prices artificially high and this has stopped commercial deployment in this market. Demand for ISDN is now so high that there are not enough services available, again preventing lower prices. Current changes in the carriers' polices are too little too late. It is now up to the cable and satellite companies to fill the gap. Digital satellite TV is particularly well suited for this market, but cable TV is certainly also a big contender. With cable modem prices possibly falling under US$400, all is set for the cable TV assault on the market from 1998 onwards.

In the United States and many European markets, however, the current cable TV plants are not up to a level where they will be able to provide reliable, two-way data traffic over their networks.

While some of the cable TV networks around the world might be better placed to provide a reliable service, without a U.S. cable-modem market there may not be a sufficient mass

TABLE 7

Intelligent Network Deployment Around the Globe

Country	1	2	3	4	5	6
Australia	yes	yes	yes	yes		
Brazil	yes	yes	yes			
Canada	yes	yes	yes	yes	yes	yes
France	yes	yes	yes			
Germany	yes			yes	yes	
Hong Kong			yes			yes
Japan	yes	yes	yes	yes	yes	yes
Korea	yes	yes	yes			
New Zealand	yes	yes	yes	yes	yes	
Singapore	yes		yes	yes		
UK	yes		yes	yes	yes	
USA	yes	yes	yes	yes	yes	yes

(Source: Paul Budde Communication based on data from GPT)

to bring low-cost cable modems to the market. On the other hand, while the uptake for cable pay TV will initially be less than 20 percent, Internet over the cable might turn out to be the savior of cable TV.

Internet Set to Win First Round of Interactive Services Battle

Things are moving fast in the repositioning of the superhighways. It is becoming clearer that the digital set-top boxes and cable modems that would allow for a high-quality range of interac-tive services will be too expensive and too limited. Based on the TV concept, it is now accepted by most operators that it will be impossible in the short-to-middle term to turn interactive TV into a prof-itable operation.

At the same time, the Internet is moving so fast that it is quite possible for it to deliver the sort of interactive services that it was envisaged set-top boxes would deliver over cable TV networks. By 2001, the Internet network will be the same size as the current telephone network.

The TV and the PC are used in totally different ways, and it would be a very difficult marketing exercise to get people to use the TV for the higher level of interactive services (banking, shopping, education, information services, etc.). The TV is an entertainment box and only a low level of interactivity such as playing along with games, polling, and quizzes will be salable. A lot of such services can be financed through advertising and through pay-as-you-go charges via premium call services. You do not need a digital set-top box for this, the new enabling technology here is an advanced analogue set-tops. This equipment is more affordable, and it can be implemented in hybrid solutions.

What this all means is one development for cable TV and a different one for interactive services. It becomes clear that broadcasting technology will not be used for full interaction. This technology has missed the boat by not leap-frogging into the computer era at a much earlier stage. The race is now on for broadcasters to make the Internet available over their cable TV networks to the large num-ber of PCs in the homes. This can be done through cable modems. Price indications are already under the US$500 mark and still dropping. Because it is a computer application, it will be much easier to charge for this equipment as well as for the services.

This change in user equipment direction will, of course, also have great implications for network development. The hybrid / fiber coax (HFC) network is seen as the clear winner for broadband to the homes. New synchronous digital hierarchy (SDH-)-based networks are the ones to go into the business market. Combinations are already in use, such as the one employed by Optus in Australia. This company is leading the world in this respect, but is encountering massive problems with commercial deployment of the technology.

FSN and Cable Modems on Hold

There is a clear feeling of a hangover in the superhighway industry around the globe. During 1994–1995, operators and suppliers promised services such as video-on-demand, fully interactive shopping, and e-trading within a few years. Despite all of this hype, we stuck to our prediction that superhighways that could provide the above sort of services would not become available on a mass-market scale before 2005–2010. Many billions of dollars later, the operators and suppliers are winding down their investments, closing down pilots and tests, and shelving newly developed products and services. Some of the biggest advocates were no less than TCI (the largest cable TV operator in the United States), publisher Time Warner, and Viacom. In 1997, TCI and Time Warner both abandoned the whole lot and declared that interactive services and television were two totally different markets, and they could not be mixed and matched.

Set-top boxes and cable TV modems were finally launched in 1997. I doubt very much if they will find a large demand. Without the appropriate services why would you need such boxes? For the next five years at least, the superhighway will be driven by the Internet on PCs, not by TVs.

For the time being, the pay TV industry has the following choices:

- Provide an up-market, wide choice (200+ channels) at a premium price, aimed at a very quality- aware and highly demanding, affluent, and well-educated market. In the United States, the digital satellite TV market has been very successful in winning over this market.
- Provide a mass offering of more of the same in sport, movies, and entertainment to mass markets at low costs (US$15-US$20 per month). These services should be aimed at less affluent, less critical markets.

Unfortunately, there is very little in between on the current delivery platforms. In the short term, this might favor cable TV and free-to-air broadcasters. In the long run, however, customers will want, and will get, more choice. There is now a better chance that Internet-based rather than TV-based developments are going to be the trendsetters for the superhighway.

Alternative Infrastructures

New privatized IT infrastructures from organizations such as banks could easily include telecommunications services in their offerings. The same applies to energy companies, water authorities, and other utility organizations. Imagine the following: you will call your bank and the Interactive Voice Response (IVR) service asks you to press one for your statement, two to transfer money, three to make a phone call, four for Internet access. Your call will be routed over their network—they have access to your bank account so there will not be a bad-debt problem. Local calls could be free to customers, special discounts apply for loyal customers, etc. Watch out for interesting alliances between banks, utilities, and other large organizations and service providers.

But, despite all of this, most spectrum auctions around the globe are doing very well indeed, and more capacity has been sold to create a whole variety of new infrastructures in the mobile communications world, with applications of which we have not yet dreamed.

Global alliances such as Global One, Unisource, and others see a great future in infrastructure and are investing heavily in it.

Insatiable Need for Bandwidth

How can we make sense out all of this? First of all, I stick to earlier statements regarding the investments made by carriers cable TV and other network structures. "Shortly, I would like a permanent open Internet connection with full broadband capacity that allows me ongoing (video) access to information and entertainment around the world, based on an intelligent software program driven from my PC that indicates what I am interested in and only provides me with the info I want at a particular point in time. Very much in the way secretaries, personal assistants, and the good old-fashioned butlers worked." The bottom line is that, as in the past, carriers do not have to worry about demand. The software and content industry will create this demand—in a similar fashion to the PC and Internet developments.

Therefore my message to the carriers is build networks, build, build, build! Talk to any thirteen- year-old who has a PC, and check what they expect from interactive entertainment, and you will quickly be able to calculate the impact if only 20 percent of the population want to use the open broadband-based networks that would allow for the applications that people want. Increases in network capacity that are inconceivable to us now will be needed within a five to ten year period. However, only those companies highly specialized in building, maintaining, and managing such infrastructures will be able to operate in this market. There is no room for small players. Currently, small players are meddling in this market with switches, personal communication services (PCS) and other networks, but, apart from a few niche markets, none of the small infrastructure operators will survive. Most of them will have sold their hardware within a few years to the bigger players. Once competition in this market is properly established it will be far cheaper to buy from these larger players than to operate these networks themselves. On top of that, all the hassle of maintenance and upgrades in our fast-moving, tech-nology-driven market will be avoided.

Fixed Will Be Wireless, Wireless Will Be Fixed

Just when service providers (SPs) thought they would use the inexpensive PCS technology to build their own local-access capability, AT&T in the United States has indicated it is going to build a nationwide local-access system based on this technology. Last year, the company abandoned several of its vertical integration strategies in the area of value-added services. It clearly wants to be the prime infrastructure supplier in the United States. This is a clear message to all other carriers and SPs around the world: you must either be serious in the infrastructure market, or you must be a customer-focused SP. If you are not, the market will weed you out very quickly.

I do not agree with some commentators who see the AT&T move as the end of HFC networks, but as an effort to get into the local-access market. Voice communications will be transferred from fixed to mobile, and entertainment services will be moved from wireless (free-over-the-air) to cable-based networks in metropolitan areas.

In both situations, the 80/20 rule (80 percent mobile, 20 percent fixed, and the inverse for entertainment prevails).

Urgent Need for Wider Participation

The networks, however, must be open to content and service providers. Only these organizations will be able to make sure that the revenue stream starts flowing in. It will be impossible for the carriers to turn their networks into money makers. The carriers, however, must provide the best possible networks for such services, otherwise customers will use other in-frastructures to obtain the services they want.

The new concept of interworking and market specialization will be molded much like the Internet—a totally open system, dominated by no one, where everyone contrib-utes to the total success. Carriers build and upgrade the infrastructure;

service providers operate in local access, Web hosting and niche markets; content providers are developing new services; software developers are adding an amazing range of features to create a more user-friendly Internet; hardware suppliers are spurred into action with new modems, switches, security equipment, etc.; business and other organizations launch their products and services on the Internet; and banks provide the financial-transaction infrastructure needed and so on.

Services and Content

Convergence Creates New Demand for Services

The telecommunication sector is at the heart of a much larger industry—information and communication technology, or "infocommunications"—which was worth some US$1,370 billion in 1995. The convergence of the telecommunications sector with the computer and the broadcasting world is creating new synergies, most evident in the exponential growth of the Internet, which continues to double in size every year.

At the start of 1997, there were some 16 million host computers connected to the Internet and more than 50 million users. The significance of the Internet lies perhaps not so much in what it is, but what it will become. It can best be regarded as the prototype of a global information infrastructure that will lay the platform for the electronic commerce of the 21st century. Estimates of the value of transactions carried out via the Internet in 1996 ranged between US$1 billion and US$3 billion.

Services and Content: Keys to Success

While the infrastructure is of crucial importance, eventually the technology will be virtually ubiquitous. The premium, therefore, will be on content. It will be the service providers who will have to package this in such a way as to generate users' interest. Content and SPs will play a key role in making the superhighway viable for all parties involved. Made complacent by the virtual monopolistic control of their distribution systems, many of the major players have been slow to recognize that the age of mass media is about to end, replaced by a spectrum of niche markets.

Service and content providers can be divided into the following main categories (see *Figure 7*):

- Traditional full-motion-video content providers (movies, documentaries, news, sports)
- Advertising agencies and other producers of commercial video material
- Telecommunications and other interactive service providers, teleshopping, telebanking, EDI, EFTPOS, on-line, videotex
- Multimedia and software developers, currently mainly producing education and games

Other important groups, closely aligned with service providers, are content centers for digitizing and storage services, and packaging and marketing organizations.

Initially, the emphasis will be on traditional pay TV services (one-way video, movies, sport, etc.). This market is heavily dominated by the eight big studios in Hollywood. With an output of over 1,000 movies a year, the content market in India is larger than that of the United States; the Chinese market is also very large. Companies such as STAR TV (News Ltd), PBL (Kerry Packer, Australia), and MTV (Viacom) are establishing joint ventures and new alliances in these two important language and culture markets.

As the technology battles in relation to multipoint distribution systems (MDS), satellite and cable are coming to a close, the focus will now shift to content. Obtaining the rights to movies is an expensive hobby—a blockbuster costs anywhere between US$500,000 and US$700,000. With more groups com-peting for these products, prices for top products will be high and it is not difficult to get the best products but go broke before you put them on your pay TV service.

Killer Applications

There will not be one killer application that will suddenly turn mass markets on. Numerous niche market applications,

FIGURE 7

Superhighway Offerings

- enhanced analogue cable television;
- video on demand;
- video games;
- interactive video shopping;
- interactive gambling;
- distance learning;
- driver's licence renewal or tag registration;
- video conferencing;
- access to long distance telephone numbers;
- personal communication services;
- full motion video picture phone service on full sized TV screens;
- medical imaging;
- high speed data transport for businesses.

TABLE 8

Killer Applications Revenues

Killer applications	Revenues 2005 US$ billion
Information services	19
Gambling and porn	22
Advertising	34
Entertainment	55

(Source: Paul Budde Communication based on data from Ovum)

packaged in the right way, will be the way to go. The software that will allow customers to find the services and information they want will play a key role in this process. According to research information from Ovum Ltd in the United Kingdom, the major applications on the global superhighway market will be worth some US$130 billion by the year 2005 (see *Table 8*).

Cable Telephony

The first cable networks were installed in Britain between 1990 and 1992. No technical arrangements were originally made to facilitate telephony. Now, all parties involved in the British cable TV business agree that cable without telephony will never be a money maker. Telephony not only makes it possible to get a return on investment quicker than was originally anticipated, it also makes the difference between a profitable cable TV network and an average money earner.

On-line Gambling

While Australia has the dubious honor of being one of the leaders in the gambling world, other Western and Asian markets are not far behind (see *Table 9*). Regulation will be the only thing that can stop this application from becoming a killer application.

Gambling from home will take away business from casinos, hotels, and clubs. Key applications for interactive TV are horse racing, football, and other sports. Viewers will be able to place bets using a remote control, smart cards, and PIN numbers.

The gambling market will stimulate growth in the interactive market in groups aged between thirty and forty-five years. Promoters of on-line gambling argue that it will eliminate the alcohol abuse, peer pressure, and overspending that often accompany gambling in casinos and clubs.

Adult entertainment

Adult entertainment is still a very contentious issue, with heavy censorship rules in place, or on the way, in the United States, Australia, and Asia. However, conservative America is now openly presenting adult entertainment as a killer application on cable TV. After the 1980s when such services were banned wherever possible and where instead serv-ices such as homebanking and homeshopping were pitched heavily to the consumer, adult services are now back on the agenda. In the past, adult entertainment has been the driving force behind audiotex and videotex. At a certain stage, Telstra in Australia even had a joint venture running with Penthouse on their on-line service.

However, the Internet phenomenon is stirring the debate again. Rather than crucifying such services, the industry and the governments should look for proper regulation rather than prohibition. A sensible approach to such services could restrict unwanted material such as violence, sex, and child pornography and avoid usage by minors (see *Figure 8*).

Apart from this, adult entertainment could be used among other services to get people involved in the new services. This in its turn could create a viable market with more money to spend on services that yield less or no profit (education, health care, community services). All that is needed is an adult approach to adult entertainment.

Pay TV

More of the same will not entice potential users. Because of pent-up demand, however, pay TV will stimulate a high uptake of services in the first six to eighteen months in countries such as Australia and Asian countries. But unless the operators immediately switch to a large number of new services, the initial subscribers will leave these services in droves when they have to renew their subscriptions. While the issue of how many channels users really want (potentially upwards of 500) is still hotly debated, the available evidence shows that the more choice that is provided, the higher the uptake will be. For example, you would not go to a library that only holds fifty books in stock or to a newsagent that sells only their preferred newspaper and a handful of their favorite magazines. The success of special interest groups on the Internet and other on-line services, the fragmentation of the music industry over the last few years, the proliferation of specialist magazines, and the increasing interest in art-house films, are all strong indicators of the increasing emergence of niche markets. Popular culture is no longer the un-sophisticated, homogenous mass market it was in earlier years.

More and more confirmations are coming in from initial trials that have been conducted along these lines. The major reasons for the lack of success were:

TABLE 9

Gambling Statistics

Country	Gambling per head/p.a.
Australia	US$400
Hong Kong	US$370
United States	US$170

(Source: Paul Budde Communication based on data from International Gaming & Wagering Business)

• after an initial period of high usage, movies-on-demand are only used a few times a month;

• the depth and width of other services provided, such as entertainment, shopping, hobby, etc., is far too small—most information and services are already known to the user and the extra information is of little use; and

• services delivered via an analogue system are described as old-fashioned and useless as it does not even allow rewinding and pausing in video-on-demand services, features users have been used to for more than a decade on their VCR.

With the successes of CD-ROM and on-line services, a trend that will lead to sustainable growth of new interactive services will, in the short term, be based on PC-based narrowband Internet services. Over time they will grow into fully fledged broadband networks.

News Services

News, at times when you want it and about what you are interested in, has become a key element in the current cable and pay TV developments. Viewers no longer have the time and interest to make daily appointments with their broadcaster at 6pm to watch the news. In the United States, news viewing time on the major networks went down from 60 percent in 1993 to 42 percent in 1996. Local news still scores around 65 percent, but is also down, from its 77 percent high in 1993.

This is opening up the market for specialized players, providing a premium service for which viewers are prepared to pay.

Demand Side Management

Utility companies are contemplating their possible participation in the telecommunications and cable TV industry. There are compelling reasons for them to do so as they will be able to make considerable savings and, at the same time, tap into new revenue streams. Cable TV operators are already describing the energy management services as the possible killer application that will make it commercially viable to start delivering interactive services to their subscribers.

DSM is an electronic energy information system, based on telecommunication or cable TV networks, which enables customers to keep track of their household electricity use and make adjustments via an on-screen reading on their TV or PC, thus saving them money. Eventually, all of a home's appliances could be programmed at certain use levels—depending on time of day, season, and customer lifestyles.

Energy information services (EIS) include:

• Automated meter reading

• Immediate outage detection

• Real time energy usage information—customers will know how much energy they are using, allowing them to make intelligent decisions on conservation

• Real time pricing information—customers will know their current energy costs, allowing them to manage their usage or shift their activity to a different time of the day

• Customer flexibility to schedule electrical use—customers could preset lighting appliances and thermostats to control their operation

• Electronic billing and bill payment

• Energy usage detail—customers will know how the use of an appliance, like an air conditioner or pool pump, contributes to their overall energy bill

Trends favoring energy conservation and competition in power generation are constantly pushing utilities toward DSM as a way to reduce costs, minimize investment in new power plants and improve customer service. DSM can represent a saving of approximately US$400 per household. Two-way coax, used in cable TV networks, provides the level of services that utilities need, including processing power at the user end of the network, e.g., graphical user interfaces (GUIs).

It is very important for the utilities to have a broadcast-communication system. Current telecommunication systems do not provide this. Copper cable–based telecommunication networks are, for example, not able to alert the entire community to the fact that, because of certain weather patterns, the price of electricity will go up for the next three hours. Cable or radio systems are therefore better suited for utility applications.

Cable TV is a low-cost provider of communication services, plus it provides intelligence to the home to help manage the home energy system. With cable TV, there will be a separate communication path to these homes, with data going through the cable TV system and with software supporting specific energy applications in those homes.

Utilities, telecommunication companies, and cable companies have the potential to be fierce competitors in the network infrastructure market. The question is, however, is it economically viable to build three separate fiber-coax systems to each home?

In Australia, several utilities are involved in Optus's plan to use their infrastructures. However, like all other superhighway projects, DSM also is still several years away. Projects in the United States are delayed or have been abandoned. The largest trial so far involved America's largest utility company, PG&E. Microsoft was involved in this project with a software product code-named Amazone (based on cable TV networks). This development has been abandoned, and Microsoft addressed DSM in its general fiber-optic project, code-named Iceberg. Telstra is also involved in this project, but beta products are not expected for another two to three years.

This does not mean that DSM as a concept will have to wait all that time. In 1995, the U.S. and Australian energy markets were deregulated. It will not take long for entrepreneurs to enter the market and attract the utilities market with new in-

TABLE 10

Falling Videoconferencing Costs

Year	Transmission link	Equipment
1985	US$2,000	US$250,000
1995	US$20	US$1,000
2000	US$5-US$10	US$500

(Source: Paul Budde Communication based on data from ITU)

novative products and services that will take the cream off some of their most profitable markets. This trend will be very similar to what is happening in the telecommunication market at the moment.

Videoconferencing

The top end of the videoconferencing market has been growing at a compound annual rate of 48 percent, with sales of approximately US$2.7 billion projected for 1998, up from approximately US$1.1 billion in 1995. Growth in all other segments of traditional videoconferencing is stagnant and even declined slightly after 1995. After being around for over a decade, very little progress has been made so far—this despite the fact that hourly costs dropped from US$2,000 in 1985 to US$20 in 1996 (see *Table 10*).

The real interest in videoconferencing is not in the US$100,000+ systems, but in smaller scale, reasonably priced videoconferencing systems for desktop PCs and laptops, which are ideal for smaller group meetings, small business, classrooms, even for use by extended families. Videoconferencing-enabled PCs are now available for a fraction of the price usually associated with ISDN, fiber-optic, or satellite-based videoconferencing technologies. The size of the market for community-based videoconferencing systems is growing. It is becoming more user friendly than existing communication technologies and will, over time, replace most of the current systems in use. Videoconferencing is no longer the exclusive province of large corporations with significant technical and financial resources.

The growth of ISDN and the Internet is facilitating the interest in desktop videoconferencing and simultaneous voice and data transfer services for small businesses and consumers. Not as a stand-alone application, but integrated with other services, videoconferencing will be able to regain its place.

The main barrier to the market growth of mass multimedia–networking services at present lies in the high bandwidth required for high-resolution, real-time interactivity. As fiber-optic density in national networks grows, and particularly as the penetration and cost of fiber-to-the-home connections decline, demand for interactive multimedia services should expand dramatically. The diffusion of such networks will in turn depend upon the technological development of the component industries, general macroeconomic conditions affecting demand and supply, as well as the regulatory structures adopted by governments.

Web Advertising: The Year Ahead

Karen Burka
Formerly Editorial Director, New Media & Marketing Group
Cowles / Simba Information

Though still in its nascent stages, the World Wide Web advertising market has changed dramatically since its birth in December 1994, when Coors' Zima brand launched its first banner campaign on HotWired. Once limited to low-resolution, bitmapped graphics, the enhanced graphical capability of the Web awakened an indifferent advertising community. Previously dominated by a handful of high-tech and telecommunications advertisers, the breadth of companies advertising on the Web has grown to include auto manufacturers, packaged-goods marketers, travel services providers and financial institutions. By late 1996, Web creativity had become more sophisticated—expanding beyond the banner—and advertisers began receiving tangible feedback on the effectiveness of their campaigns.

Several factors are driving the Web's increased acceptance as an advertising medium. The Internet audience continues to grow exponentially, with its demographics broadening to include a higher percentage of women—key targets for many mass-market advertisers.

Audience measurement and banner standards, though still not widespread, began to take root in 1997, providing advertisers with more credible data with which to make Web media buying decisions. Improvements in technology also allowed advertisers to better target their Web messages to the most relevant users. Inspired by the increased interest from mainstream advertisers, ad-supported Web sites fell in line to support standardization of the medium.

Despite the struggles of Prodigy and CompuServe, the online medium as a whole regained some stature among advertisers as America Online launched an aggressive push to boost ad revenue, and as non-Web-based e-mail services drew millions of users.

Finally, advertising continues to be the most viable revenue stream for Web publishers. The perception that Web-based information is—and should be—free, continues to pervade users' mindsets. As a result, the Web audience continues to resist Web-based subscription services. Fear about the security of financial data traveling over the Internet also continues to hinder the growth of electronic commerce.

The long-term viability of the Web as an advertising medium is not assured, however. Web advertising expenditures still pale in comparison to more traditional broadcast and print media. Until standardized hard data about the size and demographic composition of the Web audience become routinely available, the Web will continue to trail its mass-market brethren.

Implementation of Standards is Crucial

Standards, or the lack thereof, continue to be a key issue for Web advertising buyers and sellers—and a major stumbling block to acceptance of the medium among many national advertisers. Theoretically, Web-based advertising can provide advertisers with much more information about potential customers than conventional advertising. Web sites can track the number of people that view an ad, describe them demographically and report almost instantaneously any sales or leads that result.

However, the available measurement techniques—hits, page views, and clickthroughs—have been inconsistently defined from site to site and have been of questionable accuracy. As a result, Web advertisers have been frustrated by the type and amount of data that has been made available to them about the Web users that view their advertising messages. In 1996, advertisers also became more and more frustrated by the creative and production expenses that their Web advertising efforts entailed. Every Web site seemed to have different banner sizes, specifications, and format requirements.

Basing their 1996 ad buys largely on the sheer volume of "eyeballs" that could view their ads, most advertisers limited their Web ad spending to the few sites that were drawing high traffic levels. Those sites included the search engines, Netscape, and ESPNET SportsZone.

In order for the Web to compete with more mainstream media choices, including print and broadcast, both measurement and creative standards must take root. Advertiser demands for more targeted Web ad choices are increasing, as are their demands to run campaigns on multiple sites. Adjusting each campaign to the different creative nuances of each site is too costly and time consuming, especially when advertisers have so many other media choices. In order to justify creating separate, dedicated Web ad budgets, advertisers need concrete evidence of ad performance. They also need to be able to gauge viewer habits and demographics to better tailor their campaigns to specific markets. No advertiser can afford to fire blind anymore.

In early 1997, several leading Web sites and Web ad industry consortiums initiated the development of banner and measurement standards. Spurred by the demands of big-spending Web advertisers such as Microsoft, AT&T, and Procter & Gamble, several leading Web sites announced their support for these standards, which bodes well for the future viability of the Web as an advertising medium.

Without such standards, ad-supported Web sites will continue to rely on a small group of advertisers for revenue. But with hundreds of new ad-supported Web sites launching on a weekly basis, each site's share of the ad revenue pie will shrink, jeopardizing the long-term potential of Web advertising as a business model.

Leading Web Sites to Generate $248.2 Million In 1997

Driven by increased spending by Web advertisers, revenue for seventeen leading Web sites will total $248.2 million in 1997, up 93.4 percent from $128.3 million in 1996, according to Cowles / Simba Information estimates.

The search engines—Yahoo!, Infoseek, Excite, and Lycos, among others—that offer a broad-based audience and the highest traffic on the Web, and high-tech sites—including ZDNet, C / NET, CMP's TechWeb, HotWired, iWorld and PC World Online—will continue to lead the group. News sites such as CNN and USA Today Online will attract a growing share of the mainstream consumer market advertisers testing the Web and will make the largest jump in ad revenue in 1997. Ad revenue for CNN will increase 325.5 percent to $15.0 million, compared to $3.5 million in 1996. USA Today will grow 255.1 percent from $2.9 million to $10.4 million, according to Cowles / Simba projections. Other projections are shown in *Table 1.*

Until mid-1996, the Web was a haven for the technology elite in terms of advertisers, users, and publishers. The sites generating the highest traffic numbers—and attracting the most advertisers—were high-technology publishers such as HotWired (Wired Ventures) and ZDNet (Ziff-Davis Publishing). The vast majority of Web users were either computer enthusiasts or professionals in the information technology field.

With such a narrowly defined niche, it was primarily high-technology and telecommunications advertisers that were drawn to the Web, namely IBM, AT&T, Microsoft, and Silicon Graphics. Cowles / Simba Information estimates that AT&T and Microsoft accounted for up to 70 percent of all advertising revenue for many leading Web sites throughout 1996.

There were few mass-market consumer product advertisers on the Web, and those that ventured into the medium very specifically targeted the Web's predominantly young, male demographic. Liquor marketer Sauza Tequila, for example, has been a staple advertiser on Playboy's Web site since 1995 and was one of the first of any type of advertiser to sign a one-year contract. Auto manufacturers were another breed of advertiser that quickly adopted the Web, due to its upscale male audience.

For the most part, the Web did not give the Procter & Gambles and McDonalds of the marketing world enough rea-

TABLE 1

Web Advertising Revenue for Selected Web Sites, 1996–1997

Site	1996	1997*	Change
Yahoo!	$19,000,000	$32,000,000	68.4%
Excite	14,800,000	30,000,000	105.5%
Infoseek	14,600,000	27,600,000	86.5%
Lycos	11,000,000	21,600,000	96.4%
ZDNet	9,852,600	20,200,000	105.0%
C/NET	9,100,000	19,200,000	111.0%
Netscape	8,186,900	15,600,000	90.5%
ESPNET SportsZone	6,999,900	9,600,000	37.1%
Pathfinder	6,255,800	6,535,700	4.4%
CMP's TechWeb	5,800,000	9,408,000	62.2%
HotWired	5,800,000	9,180,000	58.3%
CNN	3,525,000	15,000,000	325.5%
USA Today	2,940,000	10,440,000	255.1%
WSJ Interactive	2,776,000	5,040,000	81.6%
iWorld	2,643,900	6,000,000	126.9%
Playboy	2,619,000	4,800,000	83.3%
PC World Online	2,430,800	6,000,000	146.8%
Total	**$128,329,900**	**$248,203,700**	**93.4%**

*Represents projections

Source: Cowles/Simba Information.

son to redirect advertising dollars, outside of an experimental campaign that promoted basic corporate information. The Web audience was insular, the medium unproven and creatively limited.

By late 1996, however, the landscape began to change, prompted by shifting demographics, the continued media spotlight on the Web, and peer pressure from competitors integrating their Web site URLs into every print and broadcast campaign launched. First-time Web advertisers such as Sears came to the Web in droves, while early adopters boosted their Web budgets considerably. The most promising sign for the medium is the growing number of brand-name advertisers that have dedicated budgets for Web advertising, rather than drawing from existing research and development or traditional media coffers.

Common Characteristics of Web Advertising

Web advertising changed significantly from 1995 to 1997. More and more advertisers began reaching beyond static banner advertising to Java-based moving images, and are currently taking advantage of improved targeting technology to deliver relevant messages to more narrowly defined audience niches.

For example, in 1996, many banners were standalone marketing tactics, with no other links or jumps to more in-depth marketing information. In 1997, advertisers and their agencies widely prescribed to the theory that banners need to drive traffic to a more in-depth form of promotion, such as a home page or microsite. Cereal maker Kellogg (www.kelloggs.com) limited its Web presence to a site providing corporate information, history, and nutritional tips. Visitors could also order a branded T-shirt through an 800 number. In March 1997, Kellogg began running its first series of brand ads on The Station (www.sony.com), Sony's gaming site. The banners push viewers to a branded minisite for Kellogg's Corn Pops.

Part of what is driving the change in the Web advertising landscape is the fact that advertisers have a clearer sense of their Web goals. In 1995, advertisers launched Web sites for no other reason than seeing their competitors do it. By early 1997, however, more advertisers knew what they wanted their campaigns to achieve, whether that was building brand, driving leads or generating sales. Relative characteristics of Web advertising are shown in *Table 2.*

Characteristics of Effective Web Advertising

With the strategy in place, advertisers could focus on more effective implementation to achieve their goals in 1997. Many advertisers test different types of creative strategies, relying on month-to-month traffic reports and user sales and activity to gauge the effectiveness of their Web advertising efforts.

Arguably the most important aspect of a successful Web campaign is to have dynamic content that is updated continuously. One of the unique characteristics of the Web as an advertising medium is its ability to change "on the fly." The average lifespan for a banner ad is about 12 days. Banners will only elicit responses from first-time viewers. While the search engines can offer a significant number of new "eyeballs" day after day, the majority of Web sites do not. The characteristics of effective Web advertising are shown in *Table 3.*

Direct response techniques, such as "teasers" that pique user curiosity, are becoming increasingly effective on the Web. Banners featuring prompts such as "click here" receive up to twice as many responses as banners without such prompts.

TABLE 2

Common Characteristics of Web Advertising, 1995 vs. 1997

1995	1997
Small group of high-tech, telecom advertisers	Diverse group of consumer and business market advertisers
A variety of marketing goals combining broadcast, sales promotion, public relations and direct marketing tactics	Clearer delineation between brand marketers and direct marketers
Short-term advertising contracts, usually between one and three months in length	Longer-term contracts, particularly for keyword ad buys on search engines
High incidence of barter and value-added arrangements between sites and advertisers	Less complementary ad space given to advertisers

Source: Cowles/Simba Information.

TABLE 3

Characteristics of Effective Web Advertising

- Frequently refreshed ad creative; average lifespan of a banner is 12 days
- Prompts users to respond with commands
- Appeals to viewers' sense of community
- Attracts viewer attention with movement through Java or other animation techniques
- Convenient, easy to use; banners near navigation bar draw higher clickthrough than top-of page placement

Source: Cowles/Simba Information.

Technology is playing an increasing roll in drawing eyeballs to Web ads. Affordable technology and a more techno-savvy consumer audience are the primary reasons. Whereas a year ago, high-tech graphics were a liability rather than an asset in Web advertising, Java-based ads are now helping to boost clickthrough.

However, just as an increasing number of Web advertisers are using direct marketing strategies and tactics, Web ad response rates still largely mirror direct response rates. The average clickthrough rate has increased from around 3 percent in 1996 to 5 percent or 6 percent in 1997, but that can mostly be attributed to higher Web traffic rates.

Advertisers are finding that they need to provide value-added or community-oriented content and services packaged around their advertising. America Online's custom advertising packages allow advertisers to create their own centralized mini-marketing forum within the on-line service. This package has historically been the company's most popular ad program—and one of its most expensive. Pricing ranges from $60,000 to $200,000. Oldsmobile conducts its interactive marketing exclusively through the America Online customized program, with the exception of a Web site at the General Motors corporate home page.

Advertisers Expand Beyond Banners

As the number of advertisers on the Web has increased, so has the competition for viewers. Advertisers are recognizing that banners are not enough to grab viewer attention. Many advertising campaigns in late 1996 and early 1997 included sponsorships or content-oriented promotion, rather than banners. Web sites are cooperating with advertisers to develop these types of ad vehicles—and charging premium prices to do so.

For example, The Station offers two advertising packages in addition to a banner ad program that are priced up to $75,000 per month. The first package, called entitlements, lets advertisers sponsor third-party content. Pontiac currently participates in the program, sponsoring the Jeopardy! Online College Challenge, in which it uses a combination of targeted marketing, prizes, and content integration to further the Pontiac brand. The Station also offers packages in which marketers sponsor content, such as Sears running a "What's The Weather?" game. Children and parents play the game together; users can register for Sears's Kid Vantage frequent purchase program.

Interstitial ads, which are full-page ads that pop up onscreen for as long as three minutes as Web pages download, are also being used by sites such as HotWired and PC World Online. Interstitial ad rates average a CPM of $150.

Web/Online Ad Rate Strategies

Ad-supported Web sites and on-line services continue to struggle with a variety of rate models that value the medium's advantages over traditional media, while downplaying the lack of audience measurement standards and the inability to pinpoint the exact number of unique users that interact with each advertiser's ad.

In its earliest stages, Web advertising rates were time-based and went as high as the market would bear. In October 1994, for example, HotWired set Web ad rates at $30,000 for 12 weeks, which quickly became the standard for many Web sites. Other sites followed HotWired's lead, pricing ads within a comfortable range of its $30,000 per 12-week period. In October 1995 when Conde Nast's CondeNet launched its Epicurious Web service, ad rates ranged from $24,000 to $30,000 for a three-month sponsorship.

With little proof of audience reach or effectiveness, though, many advertisers balked at such rates. In addition, Web sites were more than willing to go off their rate cards to attract advertisers. Sites with print or broadcast counterparts, such as cable television channel Arts & Entertainment, threw in Web ads as added value for a larger ad commitment to their core media properties.

Ad Rates Reflect Guaranteed Usage

In late 1995 and 1996, many leading Web sites restructured their rate cards to reflect the different audience demographics and usage volume generated on each Web page. They also started basing ad rates on guaranteed usage. Initially, Web sites guaranteed advertisers a minimum number of impressions or page views within a prescribed time period. Ad buys were evaluated on a price-per-impression (PPI) model (also referred to as cost-per-impression, or CPI), which indicated how much it cost to make one impression. PPIs among leading sites ranged from less than $0.02 to $0.20.

The PPI model changed the Web advertising paradigm from time-based to volume-based, more closely mimicking traditional media buys. However, it also created spot or remnant ad space for high-traffic sites that met their usage guarantees before the allotted time period. As a result, by mid-1996, most high-traffic sites went to a strictly usage-based ad rate model. Several PPIs are listed in *Table 4*.

Advertisers Seek Clickthrough Rates

In their quest to more accurately measure the performance of Web ad buys, leading Web advertisers—namely Procter & Gamble, Microsoft, and AT&T—in mid 1996 and early 1997 began requesting clickthrough-based ad rates. Several ad rates are listed in *Table 5*. Clickthroughs take the impressions-based model a step further to quantify the number of users that actually click on the advertiser's banner to view more in-depth marketing information that resides on either the advertiser's or host's site. In 1996, clickthrough rates averaged 2 percent to 3 percent—mirroring the direct marketing industry's average response rates. Clickthrough rates rose throughout the year, the result of higher traffic numbers and more effective banner creation and placement.

For example, Playboy's site claims a 2 percent to 9 percent clickthrough average. Advertisers using Java- based banners and push technology are drawing on the higher end of that average. Time Inc. New Media's Pathfinder is generating 2 percent to 5 percent clickthroughs, depending upon ad placement. Some sites, such as Web directory Switchboard, claim clickthrough rates of 11 percent or higher.

Clickthrough rates level the playing field for niche sites and direct marketers on the Web, because they are based on the

TABLE 4

Price-Per-Impression (PPI) for Selected Leading Web Sites

Site	Ad Buy	Rate	Guaranteed Impressions	PPI
Yahoo!	General rotation	$20-$60 CPM	1 million	$0.02-$0.06
HotWired	468-by-60-pixel banner	$80 CPM	100,000	$0.08
Pathfinder	Run-of-system rotation	$20-$50 CPM	1 million	$0.02-$0.05
C/NET	General rotation	$15,000 per month	200,000	$0.07
Infoseek	Run-of-site	$18 CPM	3 million	$0.018
Playboy	Banner	$15 CPM-$12 CPM	50,000	$0.012
PC World Online	468-by-60 pixel banners	$70 CPM	NA	$0.07
iWorld	Banner rotation	$100 CPM-$200 CPM	NA	$0.10-$0.20

Note: Rates calculated to reflect gross as of 3/31/97.
NA = Not applicable.

Source: Cowles/Simba Information.

TABLE 5
Ad Rates for Selected Web Sites

Ad Rates For Selected Web Sites

Site	URL	Ad Buy	Rate
Excite	www.excite.com	General rotation	$20 to $60 CPM; 1 million guaranteed impressions
ZDNet	www.zdnet.com	Home page	$35 to $80 CPM
ESPNET SportsZone	www.espn.com	Home page	$100,000/Quarter
USA Today	www.usatoday.com	General rotation	$30 CPM (no guarantee)

Note: Ad rates as of 3/31/97.

Source: Cowles/Simba Information.

TABLE 6
Price-Per-Impression (PPI) for Selected Leading Web Sites

Selected Web Site CPM Rates

Site	URL	Ad Buy	CPM
SportsLine USA	www.cbs.sportsline.com	General rotation	$20-$50
WhoWhere	www.whowhere.com	Search page	$10-$15
America Online	www.aol.com	Channel placement	$30-$60
PointCast	www.pointcast.com	Channel placement	$40
BigBook	www.bigbook.com	Search page	$30-$200
CompuServe	www.compuserve.com	Channel placement	$35-$47
Infoseek	www.infoseek.com	Search page	$50
Quote.com	www.quote.com	Channel placement	$30-$35

Source: Cowles/Simba Information

quality of prospective customers, rather than the quantity. A site that attracts only 10,000 visitors per week can compete with Infoseek or Netscape for advertisers because of a higher clickthrough rate.

However, many sites are resisting clickthrough-based rates because such rates hold the Web site host accountable for the effectiveness of the advertiser's banner creative and placement. Many Web sites believe that their responsibility is to drive traffic to the advertiser's banner and that it is up to the advertiser to create a compelling ad banner that prompts users to click through to more in-depth marketing information.

Because of these issues, many sites continue to look to traditional media paradigms for an ad rate model. In 1997, cost-per-thousand-based (CPM) rates dominated the Web advertising market. Several of these Web site rates are shown in *Table 6*. CPM rates favor high-volume sites, which control the lion's share of the Web ad revenue pie. The Web advertising community's reliance on a model that touts quantity over quality will hurt the medium in the long run because it discounts the interactivity and immediacy of the medium—the very advantages that Web advertising offers. As long as accurate audience measurement standards continue to elude Web sites and advertisers, though, both will continue to base their Web ad rates and buying decisions on a model that feels comfortable and safe.

It is unlikely that an ad rate standard will emerge on the Web in the next 12 to 18 months, given the variety of models currently being used. Advertising on the Web will represent a hybrid of various marketing programs and goals. CPM-based rates will work for broad-based, image advertising campaigns, while a clickthrough model will be more appropriate for advertisers looking to generate leads or sales from their Web ad messages.

On-Line Banking Technology: Status and Security Issues

H. Michael Chung
Department of Information Systems
California State University, Long Beach

Abstract

Delivery of a value-added financial service is critically dependent on the use of computer and communications technologies. As seen in the growing competition in the banking industry, the imperative for better pricing and customer service is greater than ever before. This paper addresses the current status of on-line banking technology issues, focusing on the telecommunication security enhancement that could make things better for the banking business and its consumers.

Changes in Banking Industry

Changes in information technologies, the regulatory environment, and industry competition have deeply affected the ways in which banks deliver their services, many of which resulted from innovations in computer and communication technologies and their applications. These include automated teller machines (ATMs), telephone banking, Internet-based banking, smart cards, and electronic cash, among others.

In addition to technological developments, the impact of regulatory forces was significant. For example, the Riegle-Neal Interstate Banking and Branching Efficiency Act of 1994 removed many restrictions set by the Glass-Steagall Act of 1933. Competitive forces influenced the financial services industry as well. In addition, deposit and loan functions are becoming more specialized, and / or being performed by non-traditional financial services industries.

On-line Banking

Status of On-line Banking

About 2.5 million U.S. customers currently do banking online. In United States, the number of households using on-line banking services is projected to rise to 18 million by the end of 2002 (Denton, 1997). The American Bankers Association expects that home banking will grow dramatically by 1998, but will still constitute a small fraction of total bank transactions. (See *Figure 1* for the projection for on-line banking and *Figure 2* for banking transactions by methods.) Citibank reports that 400,000 of their customers subscribe to on-line banking. Some banks offer on-line applications for credit cards, home-equity loans, and consumer loans. Customers can customize products and services, altering terms for loans in order to accommodate their individual circumstances. Such growth was facilitated by the penetration of affordable PC's with a variety of applications and faster and economic Internet access and on-line services. In addition, a number of banks began to offer their electronic services for free at least for a limited time period. Such growth will apparently continue given the aforementioned environmental conditions.

The value of on-line banking for the banks includes the ability to attract new customers, to retain existing customers, and to market new services. On-line banking also offers the potential to dramatically decrease the cost of processing banking transactions. Surveying the role of information technology in the banking industry, Chung and Bhoopatraju (1994) report that cost reduction is the primary goal (see *Table 1*). Another survey shows that the average costs of on-line banking transactions is approximately only 10 percent of full service banking at a physical branch (Booz Allen and Hamilton, 1997). In addition, on-line banking could reduce the attrition rate by allowing customers to continue accessing the bank's services even after they move (Bamford, 1997).

On-line banking services include looking up an account balances, transferring funds, getting statement data, and paying bills. Additionally, stock quotes and investment services are often offered. A large number of American banks integrate on-line banking with Quicken or Money personal-finance software packages, while some financial institutions use proprietary software systems for access. A uniform software standard has not yet been developed.

Issues in On-line Banking

An important factor for the success of on-line banking is consumer acceptance. Users will not embrace on-line banking until it becomes as simple and comfortable as getting cash from an ATM. While ATMs have been around for over 20 years, about 30% of consumer banking transactions are done through ATMs these days. While Internet users may be well suited for working with a desktop computer technology, it may take a massive shift in consumer behavior of those other than the early adopters to get used to the idea of on-line banking.

For Internet users to smoothly adopt on-line banking, the online banking services need to be attractive enough to switch from the ways they currently bank. For example, the concept of getting cash via on-line banking should be addressed. More importantly, consumer awareness needs to be im-

FIGURE 1

Projected Growth Rates for Electronic Banking (Millions of Households; Jupiter Communications, 1996)

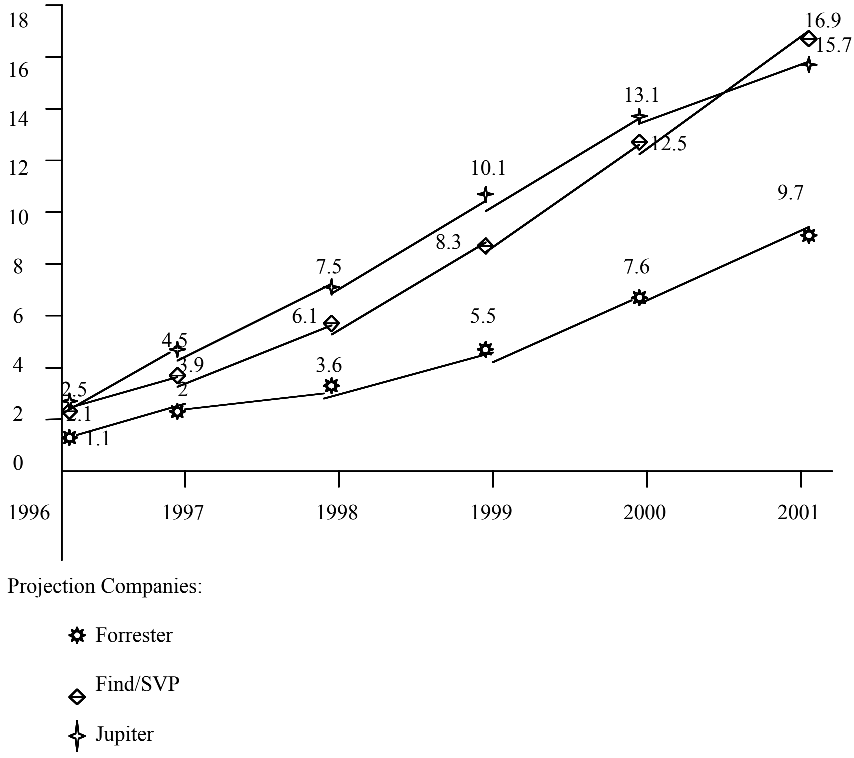

FIGURE 2

Transactions by Method

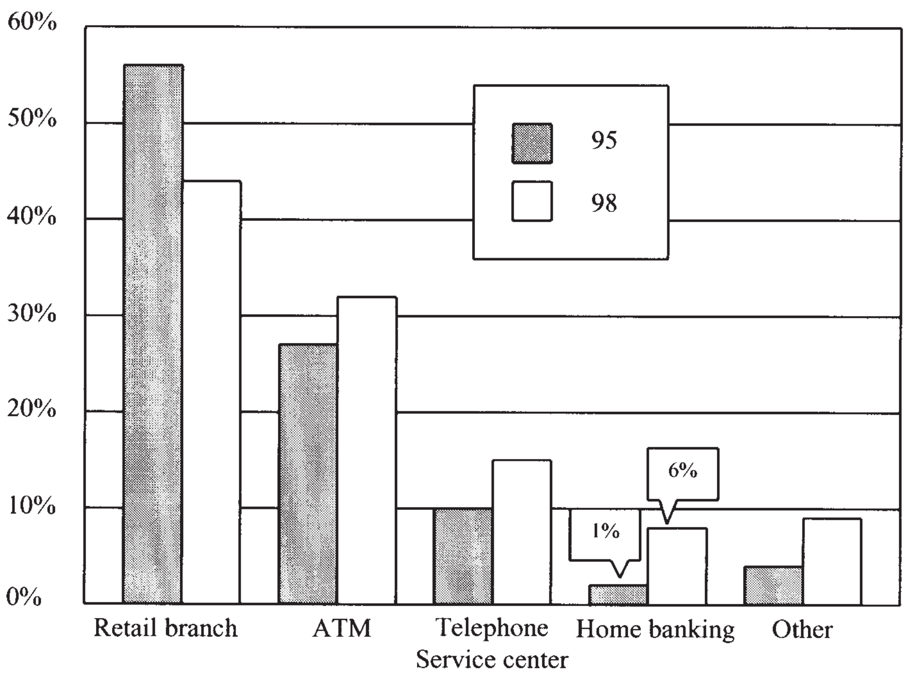

TABLE 1

Role of Information Systems in Banks

Purpose	Number of Banks
Cost reduction	9
Quick Access of Information	8
Fast Processing	7
Improved Customer Service	6
Offering New Products	6

proved. On-line banking services have not been marketed well. Also, usage fee is a concern. The cost of hosting Web sites or contracting out to an Internet service provider may not look significant compared to the benefits the service can draw. However, operating and managing the system in a competitive industry is not a trivial issue. Furthermore, interfacing electronic check payment with the traditional way of wiring checks results in a significant delay especially for those who have higher expectations on on-line banking (Bamford, 1997).

Security is an another major issue in on-line banking. Typical security measures considered are encryption, authentication, and a firewall. Encryption is the transformation of data into a form unreadable by anyone without a secret decryption key. Authentication is the use of digital signatures. A firewall is a key-system or a group of systems that enforce an access control policy between networks. The value of convenience should also be taken into account against tightening security, though. Each measure taken to increase the security of a system may make the system more difficult to access.

Finally, consumer confidence is a big issue in regards to the protection of on-line consumer transactions against third-party threats. Consumers that are unsure of security are unwilling to provide payment information over the Internet. The next section describes the security issues in detail.

On-line Banking Security

Encryption

There are two kinds of encryption: secret-key and the public-key. In secret key cryptography, a secret key is shared between communicating peers and is used to encrypt and decrypt messages on either side. All parties must know and trust each other, and have in their possession a protected copy of the key.

A public-key cryptosystem is a relation where the two keys are mathematically related so that data encrypted with one key can be decrypted using the other. A widely used public key cryptosystem is RSA by Ron Rivest, Adi Shamir, and Len Adleman. In the RSA method, each participant creates two unique keys, a "public key," which is published in a public directory, and a "private key," that is kept secret.

Data that travels through cyberspace during an on-line banking session, from a password to an instruction to pay a bill, becomes a string of unrecognizable numbers before entering the Internet. Both the bank's computer and the browser a consumer uses to surf the Web understand the algorithms that serve as locks on the doors of the consumer account information. Each time on-line banking session begins, a user computer and the bank's system agree on a random number that serves as the key for the rest of the conversation.

Authentication

Authentication refers to the process of verifying the claimed identity of a principal. There are two different authentication approaches: password-based and address-based authentication. In password-based authentication, protection of resources is usually achieved by direct log in to each host accessed using passwords, with users selecting passwords and transmitting them. There are some drawbacks to this type of authentication, such as the fact that users tend to select passwords that are easy to remember. The transmission of passwords is exposed to passive eavesdropping and subsequent replay attacks (Davis, 1998).

One way to overcome the problems of password-based authentication is address-based authentication. Address-based authentication does not rely on sending passwords around the network, but rather assumes that the identity of the source can be inferred based on the network address (Davis, 1998).

Secure Socket Layer

Secure socket layer (SSL) addresses both authentication and encryption. Authentication ensures that both the user and server are who they say they are. From a user's standpoint, this usually means little more than entering user identification and a password. But authentication is more than an old-fashioned terminal-emulation user log-on by scrambling the user id and the password information before transmission.

There are three basic properties in SSL (Oppliger, 1998). First, the communicating peers can authenticate each other using the public key cryptography. Second, the confidentiality of the transmitted data is protected, as the connection is transparently encrypted after an initial handshake and session key determination. Third, the integrity of the transmitted data is protected, as messages are authenticated and checked for their integrity.

Digital Certificates

Digital certificates verify the identity of a sender, place a secure seal on a message, and provide proof that a transaction has occurred. The system verifies a digital-certificate user's identity offline. Then the information is put into the authority server, which generates a public key and issues a digital certificate to the user, along with a related private key. The user can then encrypt data using the private key and a recipient's public key. The recipient then decrypts the data with the private key and the sender's public key-verifying the identity of the sender in the process (Bruno, 1998).

The digital-certificate standard currently in use is X.509. It supports SSL encryption and other encryption schemes. Both Netscape Communicator and Microsoft Internet Explorer support the X.509 certificates. However, financial institutions that want tighter security may opt to issue digital certificates to on-line banking customers.

Certificate authority (CA) is based on public key cryptography and typically has three components: a database that stores public keys; a cryptographic engine that generates the actual certificates; and a public key infrastructure engine that tracks the expiration date of issued certificates (Bruno, 1998). CA is built on a hierarchical or a cross certification model. In the hierarchical model, a central CA is deployed in a secure vault, while satellite CAs are kept under lock and key at different business units. On the other hand, the central CA validates all the satellite CAs. They in turn issue certificates to end-users on the network. Thus, all certificates can be traced back to the root certificates. The central CA is hard to compromise while it represents a single point of failure. However, the hierarchical CA may not map very well to some business models because of its nature of top down relationships.

The cross validation model shares trust among CA's that are known to one another. In this case, the digital certificate has already been signed by trusted CA's and serves as a letter of introduction. Unlike the hierarchical schemes, shared trust schemes do not represent a single point of failure. However, there is no central CA to automatically update all CA's if there is a security breach.

The primary issues for CA are scalability and performance. In this regard, the hierarchical model may be more scalable. While the performance of CA over the network is more of a hardware issue, it bears a significant impact on consumer behavior.

Secure Electronic Transactions (SET)

Another type of certificate is secure electronic transactions (SET), which is rapidly being deployed for financial transactions. SET is intended to provide encryption and authentication for electronic business transactions. It achieves it by establishing a system of encrypted communication using public and private keys, digital signatures, and authenticating certificates.

While many generally agree that SET's authentication system is effective, there are problems SET does not address (Haight, 1997). The specification says little about how banks approve customers to receive certificates. It doesn't guarantee "nonrepudiation." If a user has a SET certificate and software on a computer at work, another person could use it with a password-cracking program to guess the password. One of the advantages of SET's use of identifying certificates and keys is that the merchant's computer do not possesses the cardholder's actual credit card number during the transaction. Many merchants, however, use customers' credit card numbers to index various records in their legacy computer systems. The SET specification, therefore, allows banks to give credit card numbers back to merchants after the transaction. The result is that the numbers may be stored on the merchants' computers and, potentially be vulnerable to attack.

Firewall

A firewall is a collection of components placed between two networks that collectively have the following properties (Oppliger, 1998):

- All traffic from inside to outside, and vice versa, must pass through the firewall.
- Only authorized traffic, as defined by the local security policy, will be allowed to pass.
- The firewall itself is immune to penetration.

When a private network is connected to a public network, an intermediate system can be plugged between the private network and the public network to have a security wall. The aim of this intermediate system is to protect the private network from network-based attacks that may originate from the outside world, and to provide a single checkpoint where security can be imposed. By setting the firewall, all traffic in and out of the private network can be enforced to pass through this single, narrow checkpoint. This checkpoint also provides a good place to collect information about system and network use and misuse.

Packet filtering and stateful inspection are common technologies used in a firewall (Sample, 1998). Packet filters examine the details of incoming Internet protocol (IP) packets, and then either accept or drop them based on predefined rules. These acceptance and denial decisions are based on several criteria, including source and destination IP address, upper level source and destination port numbers. Packet filters implemented on routers examine packets at the network layer (layer 3) of the open systems interconnection (OSI) model, and are therefore independent of applications. Because they are not tied to applications, packet filters provide good performance. The downside is that they cannot comprehend the nature of a particular communication session and are vulnerable to hackers who may try an application-level attack. A packet filter examines only the packet header and does not examine the data within the packet.

Another type of the firewall technology is stateful inspection, which intelligently looks at sessions occurring between client and server. Stateful-inspection firewalls intercept packets at the network layer and then make a security decision based on information obtained from higher layers. Data can be classified as "good" (data that the security policy permits to be passed), "bad" (data that should not be passed), or "unknown" (for which no policy has been defined). In the stateful-inspection technology, data known to be good passes through, data known to be bad is denied, and data that is unknown simply passes through, since at this point, the firewall software acts as a packet filter.

TABLE 2
Business and Information Technology

- Reflecting the strategic vision of the business
- Seeking the top management support
- Visualizing new delivery mechanism from the perspective of the users
- Monitoring the customer satisfaction
- Communicating the benefits of the system to the users

Firewall Policy

There are two levels of the policy that directly influence the design, installation, and use of a firewall system. The higher-level policy, the service access policy, defines the transaction control protocol / Internet protocol (TCP / IP) and services that should be allowed or denied from the protected network, how these services should be used, and how exceptions to this policy are handled. The lower-level policy, the firewall design policy, describes how the firewall actually goes about restricting access and filtering the TCP / IP protocols and services according to the services access policy (Oppliger, 1998).

Conclusions

On-line banking will continue to grow in importance as more businesses go on-line and Internet users mature from early adopters to the mass market. The primary gain for both consumers and banks would be cost reduction. However, the Internet-based users will have to go through a change in the ways they use the banks. Banks can facilitate this change by ensuring the system's security for consumers and by offering a user-friendly interface.

Banks will also have to go through a behavioral change in the ways they approach their customers. Simply moving their customer service programs on-line will not be effective. A smart card that allows consumers to electronically withdraw cash from a bank account from home may be offered with on-line banking to address the ATM cash issue. However, enormous amount of capital needed to deploy new technology should not be overlooked. Management may lose their track if they are infatuated with the technology itself or merely following the fad of an on-line system without considering its business implications. In this regard, it is critical and still valid to review the Marr et al.'s (1990) suggestion for the business aspect of information technology (see *Table 2*).

Security is a major issue in on-line banking: encryption, authentication, and a firewall should be carefully considered. However, the value of convenience should also be taken into account against tightening security. Each measure taken to increase the security of a system may make the system more difficult to access. Most importantly, consumer confidence in regards to the protection of on-line consumer transactions against third-party threats is a critical issue for the banks to communicate with and convince the customers.

References

Bamford, R., "Electronic Banking," Working Paper, Stanford University, 1997.

Booz Allen and Hamilton, unpublished report, 1997.

Bruno, L., "Certificate Authorities," *Data Communications,* March 1998.

Chung, H.M. and V. Bhoopatrajoo, "Impact of a Quick Response System in Banking," *International Journal of Computer and Engineering Management*, Vol.2, No.3, pp.42-29, 1994.

Davis, B., "Certificates," *Information Week*, March 23, 1998.

Denton, N., "Net pays dividends: Online Banking," *The Financial Times*, p.13, June 25, 1997.

Haight, T., "SET," *Network World*, September 22, 1997.

Jupiter Communications, unpublished report, 1996.

Marr, N and Prendergast, G., "EFTPOS: The Perils of Cost-Driven Venture," *Services Industry Journal*, pp. 748-758, October 1990.

Oppliger, R., "Internet Security," Artech House, Boston, 1998.

Sample, C., "Firewall," *Network Magazine*, March 1998.

Singletary, M., "Players with Paperless Money: As Electronic Banking Surges, Va.'s. High-tech Firms Make the Difference," *The Washington Post*, p. F01, September 16, 1996.

Realizing Opportunities Through Wireless Technologies: Next Generation Value-Added Services and Wireless Internet Access

Brian Daly
Project Lead, Technical Team
AT&T Wireless Services

Not long ago, voice was the only service available to wireless and wireline subscribers. Over the past 20 years, however, technology has advanced and multiple networks have appeared on the horizon. Paging networks allow for one-way and two-way messaging. Through message centers and service control points, more advanced services are available through the wireless network, such as alphanumeric paging. In the wireline network, intelligent peripherals (IPs) and the advanced intelligent network (AIN) provide advanced services. E-mail, World Wide Web servers, and information services are booming on intranets and Internets. In the future, the Internet will merge into the wireless network as it currently has in the wireline network. In addition, wireless and wireline networks will converge to give subscribers value-added services. This essay will discuss the opportunities present in wireless technologies, the convergence of various telecommunications networks, and wireless Internet access.

The Coming Convergence

Wireline and wireless networks are beginning to merge, accompanied by the merger of the Internet and intranets with the wireline network. For example, browsers have made the Internet part of everyday experience. The Internet is now one of the enabling technologies for low- and high-speed data. In addition, voice and multimedia are appearing on the World Wide Web. Internet and intranet telephony has begun competing with the wireline network, something unheard of a few years ago.

Basic class services have been around for several years now, as has the integrated services digital network (ISDN) and data services. In the wireless network, many class and data services are merging. In the paging network, mobility, paging, and messaging services are also available on the wireless network. The delivery-anywhere option is one coming development. This option would allow, for example, e-mail from the Internet to be received anywhere—through a pager, cellular phone, PC connected to a wireless device, or the wireline network at home.

These advanced services are leading to a trend for the convergence of networks and services. The paging, wireline, wireless, Internet, and intranet networks will merge to provide many of the services now available through each of these networks separately. Some of these services would be available through private and residential systems, which come through the wireless network. These services include personal communications systems (PCS), packet data over the wireless network, and personal base stations that can be used as cellular phones for mobility and as cordless phones at home. Wireless local loop allows for more bandwidth availability to the home than what is available over the traditional, two-wire copper wire, wireless modems, and circuit-switched data. All of these technologies will bring about the merging of the four networks into one. Any of these services can then be offered on any of the networks.

Service convergence will make it easy for subscribers to use advanced services. Customers will have service mobility with the same look and feel from any wireline or wireless access point. In addition, users will have access to services from any location, at home or on the road. These conditions create a global market. The Internet, for example, is a global network. Carriers must think globally when addressing service mobility.

For the basic end subscriber, access to services must be simple. Most people who use cellular phones, for example, are not necessarily computer literate. Advanced services must be simple enough for the common person to use. Carriers must also meet customer demands for voice and service quality, as services must be high quality for subscribers to use them. Wireless coverage and services must encompass all areas within which subscribers may roam during the day. In addition, the cost must not be prohibitive.

Wireless carriers face many challenges, such as competition. The federal government recently auctioned off a great deal of spectrum, so many new carriers are coming into the market. Competitive digital standards and worldwide alliances are further defining this competitive landscape. Business operation challenges are a factor both in the networks that are being deployed to support digital and advanced services and in the system complexity, especially for certain technologies. The wireless intelligent network (WIN), for example, is a highly complex network to implement. Operational challenges also involve average subscriber revenue. Service provider revenue must keep up with the operational deployment cost for services. Capital efficiency is a basic business operation challenge. Marketing challenges for wireless carriers relate back to the competition: carriers must retain and recruit more customers to pay for services and find new revenue sources for existing and future operations.

Learning from the Internet

The Internet megatrend provides some lessons for network convergence. Open protocols and standards have driven the growth of the Internet. For example, hypertext markup language (HTML), hypertext transfer protocol secure (HTTP), and transmission control protocol/inter-network protocol (TCP/IP) are the basic protocols used on the Internet, and they have been standardized by the Internet community. The use of open standards, open Web servers, and browser technology allows content providers to develop services that can be used universally and quickly by all Internet users.

The Internet also allows public information to be readily obtained. The phenomenal growth of private-information intranets through corporations and e-mail have played a part in this Internet megatrend. This growth has major implications for the telecommunications and computer industries. In particular, it creates the opportunity for convergence on wireline, wireless, and Internet networks.

The wireless industry can learn a lesson from the Internet and apply it to the future of convergence and new services. The wireless industry has an opportunity to support a worldwide open standard. When broadly supported, open standards promote incredible growth. This lesson has been learned from HTML and HTTP, as supported on the Internet. Once the Internet defined open standards, phenomenal growth occurred in the industry.

Customers prefer feature versus protocol competition. Multiple Internet standards will not generate many application providers. Currently, several browsers are available on the market. Without a common standard of HTML on the Internet side—if Netscape chose HTML, for example, and Microsoft Internet Explorer chose another markup language—the vast growth of the Internet that has occurred over the past several years would not have happened. Content providers would join one camp or another, eliminating a common ground for the development of services and applications.

Another lesson learned from the Internet is that browsers must be easy to use. People have learned to navigate the Internet with browsers quickly, because they are user friendly and provide a clear user interface for advanced services. Corporate users and consumers appreciate the freedom that browsers provide. The final important lesson learned from Internet usage is that, as seen with wireless, open standards can lead to exponential growth.

Convergence or Collision?

Opportunities for convergence can be the cause of a collision. On the wireless network side, the basic architecture is the switching system. The message center, base stations, and telephones comprise the architecture for services. The message center is used for short message applications. On the intranet and Internet side, e-mail servers support e-mail applications and simple mail transfer protocol (SMTP). Web servers with HTML protocols and basic user PCs complete the architecture.

The short message service (SMS) itself will be limited to high-latency, low-bandwidth applications, regardless of the technology used. When a short message comes into a message center, the center must find the wireless mobile subscriber. This involves paging the mobile to find out if it is available for message delivery. If it is not readily available for message delivery, then the message center stores the message for later delivery. This process is a high-latency, low-bandwidth application. Low bandwidth is needed because the air-interface side typically cannot support much data rate with high-bandwidth applications over an SMS channel, regardless of technology. The need for short message service is for very small applications of a few hundred bytes, not megabyte-sized applications.

Convergence of the Internet and wireless networks raises a standards issue—what standards should define interfaces between these two networks? In the PC realm, large screens are used for viewing browser applications. The wireless network side is limited by screen size. Typically, phones are small devices with three- to four-line displays. Screen size is one factor that will limit the information that can be provided due to limited available bandwidth. Protocol efficiency is another limiting factor. It is expensive to add new radios to support data applications. The protocol that is chosen must be very efficient and tailored to the air-interface environment.

The challenge for convergence is to define open standards that will enable the creation of new services and provide a flexible architecture. Carriers want to leverage the existing investment in the wireless network infrastructure, as well as the Internet investment and infrastructure. Proven technology and standards must be incorporated into the convergence. Performance considerations require low latency over the thin pipes of the air interface, and these services must be optimized for the small screens on cellular phones. Secure and reliable service is also critical.

Thin-Client Architecture: A Solution That Works

These requirements lead to a need for thin-client architecture. A thin client is a browser in the handset that uses a standard markup language. It is also a universal client, meaning it can be used with any wireless network. With a browser in the handset, developers can upgrade applications without needing to upgrade the phone. This is an important capability because as new services roll out, the subscriber will be able to access them without having to trade in the phone or download additional software onto the phone. In addition, developers can update the applications on an application server, which makes adding new services highly cost effective. Browser technologies on the thin-client architecture provide a simple user interface and, therefore, reduced support cost.

Thin-client architecture provides benefits for carriers, subscribers, developers, and content players. Just as the Internet's exponential growth was driven by a common language that developers and content providers shared, thin-client architecture would envision the same growth for wireless. The Java programming language would fit into this thin-client architecture, but its implementation will not occur for several years. Handheld-device markup language is best tailored to cellular phone architectures at present. A standard language for browser technology that does not require software updates will allow for the rapid deployment of new services with minimal cost. Another benefit is that the entrepreneurial energy of the Internet can be reused for wireless—corporate customers with intranets can develop their own solutions internally for using this architecture. Consumers will enjoy easy-to-use browser technology and access to diverse services. Content players and software developers, using a standard language, can produce compelling information services.

Thin-client architecture offers carriers many benefits. Rapid service deployment will improve time to market. At present, developing a new service and deploying it in the field can take several years while standards are developed, equipment is manufactured, and the service is deployed. Even at the point of deployment, it may not be available to all customers, because their telephones may be older models that lack the newly available software. Browser technology provides increased flexibility, allowing services to be deployed more quickly to customers than today, while reducing dependency on vendors. More of the application development would be pushed toward content providers rather than infrastructure vendors.

Another benefit to carriers is that additional services provide additional revenue sources. New information and Internet services would increase the customer's use of cellular phones. Subscribers like cellular phones for browser and information services and, therefore, will be more apt to use them for voice services. Additionally, thin-client architecture would leverage the existing infrastructure, allowing new services to be deployed that differentiate providers.

The customer will benefit from thin-client architecture as well. Optimized protocols that are developed for handheld devices will reduce air-time expenses. For example, using typical HTML and HTTP over a wireless circuit connection would involve large air-time expenses because of the time needed to download the data. Architecture and protocols that are optimized for air interfaces will significantly reduce air-time expense by minimizing the amount of time needed to transfer the information to the phone. The browser-based application environment will reduce the cost to end users, because their phones will keep up with the latest technology. As new services are deployed, subscribers will not need to purchase a new, more advanced model or take their older phone into the store for a software upgrade.

Leveraging existing investments benefits the application and content providers, who can then use existing off-the-shelf Internet hardware and networks to provide applications. Providers can also support multiple mobile platforms. Again, the architecture would not have to be tied to the specific air-interface technology. Thin-client architecture with universally adopted standards would apply to any air-interface technology, which makes it efficient for content providers to provision it to multiple carriers. Finally, the proper choice of technology would provide robust security end-to-end for the customer and be used in certificate-based authentication procedures.

Thin-Client Architecture in Action

Figure 1 illustrates how thin-client architecture works. On the left side of the gateway is the wireless network, which includes message centers, service control points (SCPs), home location registers (HLRs), mobile switching centers (MSCs), and base stations. On the right side are the Internet and the Internet provider's equipment. The gateway with carrier applications connects the two.

If a user connected to this architecture wanted to download a stock quote, the Internet application would send information from the server into the gateway. The gateway would then deliver a notification to the telephone that it has received information.

With interim standard (IS)-136 architecture, there are two mechanisms for delivery of data information (see *Figure 2*). One is the traditional SMS approach, which would be over-the-relay data (R-Data) in the IS-136 air interface. In the other, the circuit-switched data approach of IS-130 and IS-135 is the standard used.

When the gateway receives information from the Internet, it sends a notification via the SMS approach to the phone to report that the stock quote has been received. The user's screen would then "light up" and the information would be available. The user could use interactive access through a circuit-switched data path and retrieve additional information from the Internet through a Web site.

This architecture will evolve through the introduction of packet data into wireless. Within the next five years or so, packet data will make its way into the mainstream of wireless service. Packet access supports data-intensive applications, such as personal digital assistants (PDAs), allows a laptop computer to be connected to a cellular phone, and has very extensive browsing capabilities.

FIGURE 1

Thin Client Architecture Overview

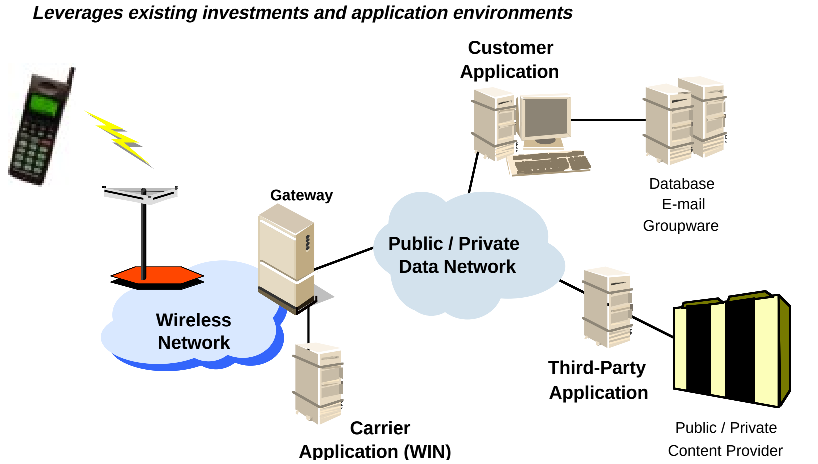

Other Applications of Thin-Client Architecture

Thin-client architecture is not limited to data services; it can support advanced voice services as well. Current standards are being developed to support WIN capabilities in the wireless network. Thin-client architecture with a browser and phone will be able to support WIN services, including typical wireline services, such as maintaining call lists, accessing on-line directories, downloading address books into the phone, controlling WIN features, and controlling basic class features such as call forwarding, call redirection, and the maintenance of screening lists.

Integrated messaging will also be possible with an integrated, two-way pager. It will provide two-way Internet e-mail and mobile-originated messaging from the phone. Eventually, this will take the short message service that is available in cellular phones or PCS today and replace it with browser technology.

Third-party companies can also develop advanced notification services. In notification and push services, information is sent from the Internet down into the phone. One example is the Pointcast application that is used on PCs. An application on the Internet can be set up to send information at certain hours of the day, or when certain transactions take place. For example, when a stock reaches a certain point, notification could be sent to the user by phone. The user then could either trade this stock on-line, make a phone call, or browse the Internet for further information. These applications are almost limitless in the Internet realm. Content providers with a standard markup language can develop almost any application they want.

Existing Web sites will be accessible through thin-client architecture. The gateway will understand that it has to interwork between Internet and wireless device applications. Therefore the information will be stripped from the HTML file and sent to the phone to avoid transmitting complex graphics or files, which requires the large browsing times that normally occur with Internet access. In the future, content providers may develop another language specifically for handheld devices.

Convergence is dependent upon open standards for common browser technology among multiple air interfaces. The fact that multiple air interfaces exist will not change. Broad acceptance of a browser technology will reduce industry risk, fragmentation, and uncertainty, causing the wireless industry to grow at a rapid pace. Competition will not be based on standards if common browser thin-client architecture is chosen.

Open standards offer nonwireless players a single, layered interface. From the Web server, information is sent to the gate-

FIGURE 2

Thin Client IS-136 Architecture

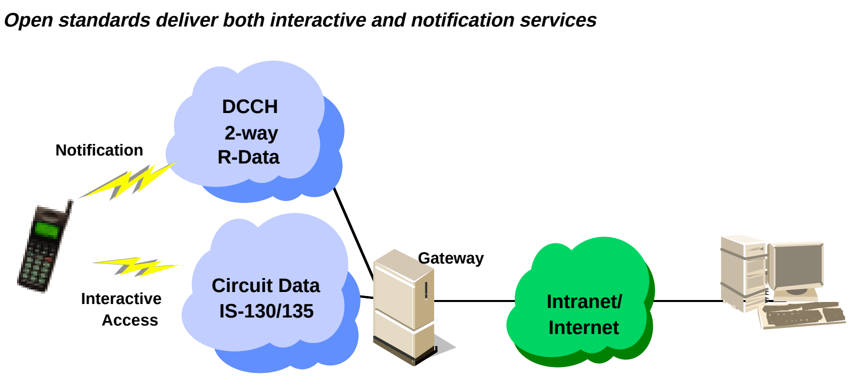

way as user datagram protocol (UDP) packets running over IP. From the gateway to the phone, UDP packets are sent via a teleservice approach in IS-136. This occurs at the relay-data level. In global system for mobile communications (GSM) or code division multiple access (CDMA), it could be imbedded within the user-data portion that SMS delivers, similar to SMS delivery to the phone. On the air-interface side, the phone is known only by its mobile identification number (MIN) or its mobile station identification (MSID). Incoming e-mail, for example, would be delivered to a particular domain name. This gateway would then translate the domain name into a phone number or MSID for delivery to the phone.

The Future of Convergence

The next generation of global standards, IMT2000, is being worked on in the International Telecommunications Union (ITU). The goals being defined for it include a high degree of worldwide commonality, compatibility of services, high service quality, worldwide seamless roaming, and the international use of small pocket terminals. In the area of services, the goal is to raise the bar, to provide voice and data services that are not currently available, to provide high bandwidth data and multimedia services, to provide bandwidth on demand, and to offer service adaptation. Thin-client architecture and browser technologies in the phone can meet many of these goals. For example, service adaptation applies to browser technology because the service is defined outside the phone and the wireless network. Some services can be defined in the Internet itself. Also, through bandwidth on demand, low-latency, low-bandwidth services can be provided over short messaging service. Higher-bandwidth services can be provided either in circuit-switched or packet mode. Thin-client architecture raises the bar to provide services that are not currently available by allowing access to the Internet through browser technology.

Any services available on the Internet today can be applied to thin-client architecture. These services are available to the cellular phone rather than to the desktop or laptop computer. Subscribers can take a handheld device with them and be able to obtain these services whenever they are within a wireless network that supports thin-client architecture.

In summary, telecommunications reform and convergence open up new opportunities. Thin-client architecture enhances voice and messaging services and provides a common platform for deploying these services to the subscriber quickly and efficiently with minimal impact to existing infrastructures. Thin-client architecture provides flexibility for increasing value-added services very rapidly. Thin-client architecture provides one way to implement these services and to meet the challenges they represent.

Call Centers and the Internet

Mary Dunlop
Vice President and Business Manager
Nortel Call Center

Businesses worldwide are increasingly using call centers to meet the changing needs of consumers. Consumers have become very accepting of different ways that businesses complete transactions. This essay will focus on the impact of the Internet in the call-center environment, beginning with business trends and technology drivers. Call-center applications and the effects of implementations will also be discussed.

Business Drivers

The trends that are driving businesses to implement call centers can be categorized into three different areas: the consumer, business impact, and technology. Consumer buying behaviors are evolving. The Internet illustrates how quickly behavior can change. Consumers are challenging businesses by entering through new digital doorways, using information appliances that range from a standard telephone to a PC Internet telephone to other multimedia devices. In the call-center business, everyone wants to know which device consumers will be using over the next several years. The answer is that they will be using all of them.

Consumer demographics are also changing. Senior citizens behave much differently than baby- boomers, who are split on accepting technology. The Generation Xers are embracing the Internet. Because of nearly worldwide Internet acceptance, a multitude of access points must be supported. In many industry segments, call centers are becoming the main model of turning revenue for business.

Access to information and communities of interest are also affecting consumers. For example, on the Internet mothers of 10-year-olds can log on and find out what fourth-graders are doing around the United States and what sneakers or baseball mitts they prefer. These communities of interest are modifying the approaches that business customers must consider when they make arrangements to accept these callers and provide one-to-one call treatment.

Businesses are being challenged as never before because they must now contend with elements of the global economy. A radical shift toward dealing with automated transactions is occurring. Right-sizing and downsizing are factors as businesses try to streamline costs. Corporations are pressured to find the most appropriate ways to right-size while minimizing the effect on the bottom line.

Right-sizing can be accommodated through the implementation of a call center, whether that center accepts calls via voice, data, e-mail, or any multimedia transaction. The call center is a means of displacing another, more costly method of doing business because it becomes the virtual storefront, eliminating real estate, the physical storefront, and to some degree, the people. While dealing with a new way of providing information, businesses are also trying to understand the implementation and the training necessary to handle the enormous amount of data arising from a multimedia consumer group. The brilliance of this implementation becomes the competitive weapon, because through that implementation businesses are grading consumers on a one-to-one basis. Information must be integrated with the host computer to allow it to respond correctly before the consumer even answers a question.

Technology has evolved to a new level. Integrating components and open standard environments are now possible. This integration presents customers with a great challenge as they examine the impact of new technology.

Much attention has been given to the Internet. Some businesses are betting everything on it, while others believe it is a passing phase. According to one author on the Internet, Nick Negraponty, “the population of the Internet itself is growing at 10 percent per month. If this rate of growth were to continue, the total number of Internet users would exceed the population of the world by year 2003.” The Internet is engaging a great deal of interest in a very specialized segment of the community. It will impact businesses through the need for multimedia interfaces for electronic mail, World Wide Web forms, face over inter-network protocol (IP), and video over IP. The infrastructure for this technology is continuing to mature as the problems caused by the interest and volume on the Internet are solved and the evolution of technology and ideas continues.

The global reach of the Internet is a critical point in the call center. The Internet knows no boundaries. The potential for new markets to reach new customers is tremendous. Mass customization and one-to-one marketing is possible, allowing businesses to reach out and provide customers with services that are tailored to their interests. Electronic commerce transactions are becoming very common, particularly in certain parts of North America. Business-to-business and consumer-to-business transactions will be primarily automated in the future. As these possibilities are realized, it becomes even more important to understand the impact of the Internet in the call center.

The convergence of all these elements offers consumers both information and new choices. In the convergence of some very key areas, a new business model is developing. The Internet is bringing the telephone companies and the Internet service providers (ISPs) together to provide circuit switching, packet switching, and cell switching. The information technology and telecommunications industries are converging. The new model that is providing solutions to the Internet and call-center environments integrates voice and data in holistic solutions that use digital switching and data networks as strong backbones to an integrated environment. Electronic commerce and call centers are coming together. The impact of the transactions and the importance of the 800-inbound calls and the World Wide Web combined results in the need to greet consumers with options. Of the 230 million people in the United States, 20 million can access the Internet. A large percentage of the 4 billion people in the world are looking for methods of accessing transactions that are not on the Internet. It is important to integrate multimedia components while offering customers the choices they need.

Technology Drivers

Whether it is the Internet, the intranet, or the extranet, these underlying technologies offer many access points and options to overlay applications based on standards, scalability, portability, rapid application development, environments, and data mining. All are becoming extremely valuable to both information-technology and call-center managers. The return on investment in a call center often occurs within days or months. The value of saving even one second is a critical component when looking at how to reduce the amount of time it takes to perform transactions. For example, a major airline reports that if it could shave a second off each call over a period of seven days across its global operations, it would save $1 million per week.

Internet access provides a return on investment in millions and billions of dollars for businesses. Much can be learned by watching other businesses engage in the implementation of the Internet and call centers. Some are engaging in decisions to integrate voice over IP call-back applications, while others are quite satisfied at just having their Web server intact and finding appropriate methods of upgrading their environments. A survey done by Nortel found that 80 percent of its business customers were interested in engaging in an Internet-enabled call center immediately. This change in attitude has occurred quickly. In 1994, a Harper's Index study found that 58 percent of adults in North America had never heard of the Internet.

The opportunities arising from the Internet are overwhelming for business. The Internet is not just global—it provides access seven days per week, 24 hours per day that is fully interactive, visual, and personal. Consumers can guide themselves to the right location. The benefits speak for themselves.

Some Internet challenges or problems are perceptions and others are realities. Momentum management, or keeping up with a continually changing world, is one challenge. Lack of security on the Internet is rapidly becoming more of a perception than a reality. Available encryption techniques reduce the security concern. The exclusiveness of the Internet can be another challenge. Internet users are a small percentage of the overall customer base. How quickly the Internet will evolve and how many consumers will adapt to it is a very valid question. To some degree the Internet is an incomplete infrastructure. Some people feel it is impersonal. The Internet, however, is becoming more interactive and more personal through new voice-interactive techniques. For some the disorganization of the Internet is problematic, but the problem of slow Internet access is a matter of network bandwidth. As bandwidth evolves, many improvements in speed will occur. In addition, some people are concerned by the perceived unreliability of the Internet, which is part of the lead environment.

Internet-enabled call centers are already being implemented. The World Wide Web is combining with the public switched telephone network (PSTN) in various ways, from e-mail environments to real-time IP voice and video. A call center that provides interactive response and is media independent can provide a myriad of services. Agents are becoming video-fax-human agents and customer-information centers.

Applications

The call-center applications that can be accommodated are as diverse as the benefits. One area that is intriguing to many businesses that want to implement the Internet is the opportunity to smooth call traffic. When dealing with hundreds of thousands of calls per week, the spectrum of interactivity can be moved to deal separately with elements that do not require an immediate response and to deal immediately with those that are real-time interactive. Speeding up traffic can be accomplished by dealing with many different components, blending human agents with data-automated agents, and having faxes and e-mail handle non-real-time delivery of information. Human agents can take live calls more quickly by off-loading much of the call traffic to an automated system. Growth of business and call volume is occurring at a phenomenal rate as consumers accept these new business methods.

The industry segments that currently lead call-center integration are financial-services companies, which generally tend to be ahead in technology arenas. Activity is also occurring in telephone companies, ISPs, and retail, travel, and health companies. Customer-service and human-resources workers are examining the need to provide one-to-one service.

The multimedia transaction center has multiple access points: the traditional phone, video kiosks, ADSI sets, and the World Wide Web. The infrastructure includes the PSTN, wide-area networks (WAN), and the Internet. Call centers can be centralized or distributed to at-home or automated agents. The availability of multimedia information anywhere, anytime, and anyplace is a critical element of providing a call center.

Differentiation through implementation is a competitive weapon that will make a difference in the future. All businesses are trying to respond to a unique set of business objectives and unique legacy investments while trying to integrate and implement the best situations to deal with their products and clients. The one quality that all call centers have in common is that they are all uncommon.

The New Electronic Commerce Infrastructure

Paul L. Evans
Director of Technology Research
A.T. Kearney

Abstract

This paper presents the results of A.T. Kearney research into electronic commerce and, specifically, the changes that electronic commerce is already triggering in the relationships and transactions among manufacturers, suppliers, distributors, retailers, and consumers.

Introduction

In general, we use the term "electronic commerce" to describe the standards, practices, and technologies that constitute a new approach to doing business. More directly, electronic commerce is the paperless exchange of business and transaction information using EDI, electronic mail, electronic bulletin boards, electronic funds transfer, and other similar technologies on private and public (e.g., World Wide Web sites) networks.

In the interconnected world of the 21st century, it will be necessary for all businesses to implement electronic commerce to meet the imperatives of an increasingly competitive marketplace. This paper describes the evolving architecture for electronic commerce technologies along with some pioneering industry initiatives and recommended actions that business executives should take to achieve a strategic advantage through the application of electronic commerce.

The History of Electronic Commerce

Electronic commerce began its business evolution over 30 years ago. In the beginning, the primary application of electronic commerce was electronic data interchange (EDI). EDI is the computer-to-computer exchange of routine business information in a standard format. Through EDI, purchase orders, quotations, invoices, shipping notices and requests for proposals, and other paper forms were successfully replaced by standard EDI transactions.

The fundamental business concept behind EDI was the use of a set of computerized forms that would automate common business transactions, saving time and money, reducing paper transactions, and streamlining business processes for the participants. These forms would be exchanged over a value-added network (VAN) connection. An EDI value-added network is a private communications network that transmits, receives, and stores messages for EDI trading partners, generally through an electronic mailbox.

In the 1970s, the transportation industry was first to develop a syntax, data dictionary, and methodology to standardize common business transactions using electronic data interchange technology. By the end of the decade, the grocery industry was next to implement EDI. Grocery manufacturers sought to establish a structure of business transactions and a common communications standard, the Uniform Communication Standard (UCS), which allowed data to be electronically exchanged among grocery industry manufactures, suppliers, wholesalers, and retailers.

In the early days, most EDI connections were limited to special purposes, often applicable to only one company or application. The format of the data usually reflected the information needs of either the sending or the receiving company. Routinely, the larger company was the one that mandated the relationship. If small suppliers, for example, wanted to do business with a big company that used EDI, they often had to join the EDI VAN or risk losing the opportunity to partner.

Over time, EDI gained acceptance within other industries (e.g., insurance and health care) but widespread adoption across industries did not happen. For the most part, only very large companies and a few government agencies had the scale, the money, and the clout to fully automate their partner and supplier relationships (that is, until now).

With the commercial birth of the Internet in the 1990s, a new generation of business networks began unfolding. The Internet (the public network), intranets (internal company Internets), and extranets (cross-enterprise connected intranets) provide a new electronic commerce infrastructure that is ultimately connectable up and down the supply chain and to consumers.

Through the Internet, the key barriers prohibiting universal electronic commerce participation (i.e., cost and connectibility) are no longer issues. Further, with Internet-EDI and new Web-based electronic commerce applications (extending well beyond EDI), today's electronic commerce offerings are maturing into a strategic set of capabilities that few companies can afford not to pursue in the future.

In the electronic commerce marketplace, Web-enabled transaction and electronic commerce software allow products to be bought, sold, and serviced via the Internet. Within the enterprise, open network software breakthroughs are coming together with intranet deployments to provide a complete suite of integrated enterprise resource planning (ERP) applications and electronic commerce products that are transforming preexisting business processes and driving competitive advantage. At the same time, bold leaders in the electronic commerce frontier are deploying extranets and forming communities (i.e., manufacturers, suppliers, distributors, retailers, and even consumers) that are strategically equipped as partners to fully integrate the supply chain and ultimately provide mass product customization for consumers.

In short, the age of electronic commerce has arrived for business. With its arrival, even the meaning of electronic commerce has been updated and modernized. Today, electronic commerce is the full collection of processes, people, technologies, and measurements that will drive strategies for competitive advantage in the digital economy. In 21st century, electronic commerce is no longer an application; it is the way of doing business.

Electronic Commerce Today and Tomorrow

A.T. Kearney's vision is that early in the new century no single or precise definition will exist to distinguish electronic commerce from business as usual. Electronic commerce will be deeply embedded in the business technology infrastructure. However, today, electronic commerce is still evolving, and because it is evolving, attempting categorization is like taking a series of photographs of the nascent cocoon at various points in time before the butterfly emerges. It is difficult to capture metamorphosis until it is complete. To date, the focus of most examination has been on the embryonic market of electronic commerce with descriptions of enabling technologies and new market opportunities—as they are emerging—not the mature specimen. We believe this kind of examination, while worthwhile, does not reveal a complete and accurate picture of the real power and potential of electronic commerce for existing companies. What's more, it might even mask or confuse strategic assessment of the critical importance of electronic commerce for today's enterprise business planning agenda.

Through our research, A.T. Kearney has discovered that immediate, significant, and powerful changes are being driven by electronic commerce in the marketplace, the internal enterprise, and the inter-enterprise business rules for existing companies. These marketplace, enterprise, and relationship changes are swift, direct, and strategically consequential.

Management needs to understand that it is not too early to prepare for these changes. In fact, those who are quick to understand and incorporate the revolutionary effects of electronic commerce, within their own enterprise and between the enterprises of themselves and their partners, are sure to provide market leadership in the new millennium.

To fully benefit from these changes and to build realistic electronic commerce strategies, business leaders must first grasp the business significance of shifts in underlying technologies that are driving electronic commerce and helping to form the new marketplace, internal enterprise, and inter-enterprise infrastructure. The emerging electronic commerce infrastructure and its related requirements will be described later in the paper.

The Coming Electronic Commerce Infrastructure

Electronic commerce involves a variety of activities, and the evolving electronic commerce technical architecture must support all aspects of these activities that range from Web transactions to trading partner interface, and from e-mail to EDI. Also, electronic commerce has substantial communications requirements. It must include a telecommunications capability that is efficient, reliable, and capable of accommodating the increasing volume of electronic commerce traffic. It must also provide gateways and network entry points that serve as high-speed telecommunication links with commercially operated VANs, as well as the interconnected networks of the public Internet, corporate intranet, and enterprise extending extranet.

The electronic commerce "virtual network" that will emerge in the new century will be an aggregate network of networks including Internet, intranet, extranet, and (in many cases) a VAN interconnection. The ultimate role of this infrastructure will be a single interconnected, interoperable, standards-based internetworking electronic commerce environment (*Figure 1*).

As shown in *Figure 1*, in the new infrastructure, electronic commerce services extend from the Internet, intranet, and extranet layers to network database and value-added network (such as VAN-EDI) deployment layers.

For the Internet layer, World Wide Web sites and electronic commerce software (e.g., transaction and electronic commerce-enabling technologies) provide lead generation, sales, customer service, order placement and tracking, external publishing, external advertising and promotional communications, and Web-based catalogs and directories for trading partners and consumers.

At the intranet layer, connectivity provides internal communications and coupled with "bridgeware" (software used to bridge the gap between new and old applications on the intranet) melds the previously disparate applications of the internal enterprise into a common, open, shared infrastructure.

The extranet layer extends the enterprise to the community of partners. Through the extranet manufacturers, suppliers, logistics operations, contractors, and outsourcing specialists are connected. Network databases and repositories ride on top of intranet and extranet deployments and provide critical, secure information to the integrated community members on a "need to know" basis.

FIGURE 1

Electronic Commerce Infrastructure

- The piece parts include: (1) the network databases; (2) business Internet, intranet, and extranet communications (open, secure and closed user groups); (3) transaction and other electronic commerce enabling technologies; and, (4) integration with private financial services and EDI VANS.

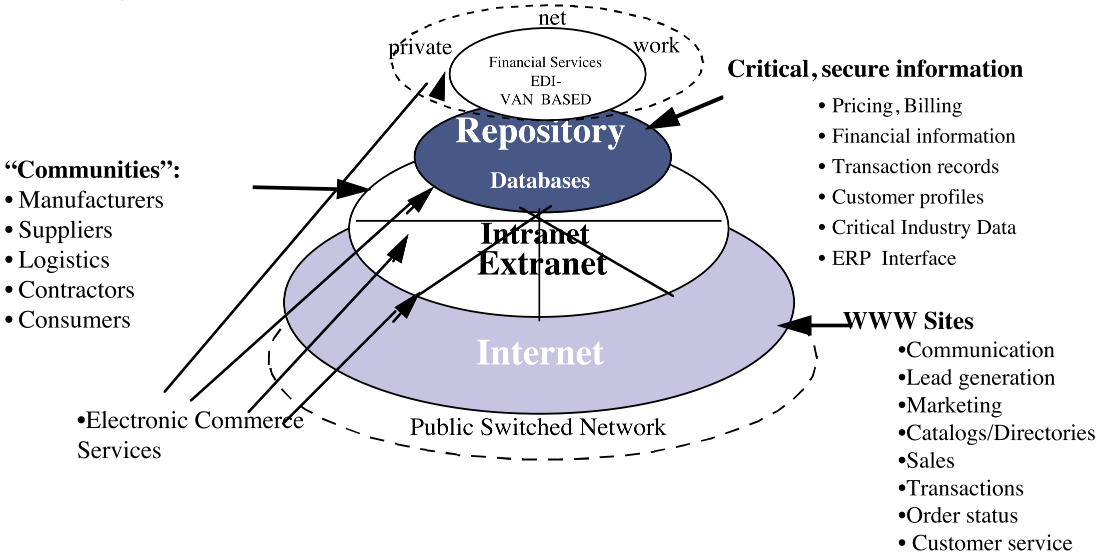

A proprietary network service layer sits atop the electronic commerce infrastructure. These network services are outsourced and accomplished over proprietary value-added networks (VANS) that interface with the virtual private networks of the electronic commerce community. Here, proprietary financial services reside along with VAN-EDI and other special services such as World Wide Web site hosting, catalog development and hosting, database programming, and management. Electronic commerce community and infrastructure management services also may have a private network component.

Enterprise Automation, Integration, and Extension

Our work shows that the execution of an electronic commerce strategy is best thought of as the result of an evolutionary technology adoption and integration process that has been underway in most businesses for quite some time (*Figure 2*).

A brief discussion of these evolutionary and essential steps toward electronic commerce is given below.

1. Internal Process Automation and Improvement

Many companies have spent the last few decades automating and perfecting their process inside the enterprise. This is the first step toward realizing the benefits of electronic commerce. Process automation is the automation of existing paper, pencil, and human tasks and processes through information technology. Of course, process automation is not new. Over the last half century, business have been engaged in some aspect of technology deployment and process automation.

Productivity applications such as word processing, spreadsheets, faxes, e-mail, electronic business forms, and, at a more advanced level, even EDI and bar coding, are examples of process automation. Recently, Web-based commerce, groupware, and Web deployment of technical support, call centers, promotional material and advertising, and employee and shareholder information over the Internet, are all examples of process automation.

Enterprise resource planning (ERP) systems such as order management, financial management, logistics management, and production management are also examples of process automation tools.

Most businesses involved in process automation have found it to be a necessary, but not sufficient factor, in gaining competitive advantage. For greater return, process automation must be teamed with business process improvement. Process improvement is the actual business process change that is enabled by the application of technologies in process automation.

The business world is not unfamiliar with this concept. In fact, it is undeniable that business process redesign and re-engineering have become the cornerstones of business strategy in the 1990s. These concepts and the ideas behind them, in their purest form, refer to the use of information technologies to reconstruct the enterprise around the capabilities of the technologies, instead of fitting the technology into existing organizational structures. Here, the important determinant of success is the recognition and execution of change in business process that is concomitant with the implementation of automation technologies.

FIGURE 2

Evolutionary Stages of Electronic Commerce

	I	II	III
Stage	**Internal Process Automation and Improvement**	**Enterprise Integration**	**Enterprise Extension Community Formation**
Sample Area of Impact	• Communications • Productivity • Operations • Functions • Customer Interface • Sales and Marketing	• Integrated Internal Systems • Enterprise Wide Computing and Communications • Integrated Applications	• Integrated External Systems • Virtual Organizations
Sample Deployments	• Internet (Public WANs) • Stand Alone Applications • Corporate LANs • Private VANS-EDI & Financial Services • ERP Applications • Business Process Re-engineering	• Intranets (Corporate WANs) • Interconnected Communications and Computing Environment • Interconnected Applications • Remote Enterprise Connect	• Extranets (Inter-Corporate WANs) • Virtual Community Constant Connect

2. Enterprise Integration

A second essential step in the evolution of electronic commerce is enterprise integration. For most companies, process automation and improvement came in single-step increments. In fact, the various technical implementations that sprouted and developed over the last 25 years most often evolved from separate family roots of production, procurement, administration, or engineering and development. Once deployed, however, these applications rarely shared data and information. By and large, until now they have remained isolated, stranded, and confined to their own technical islands with unnecessary communication bottlenecks and nonintegrated data management.

In the age of electronic commerce, unless business systems are thought of as an integrated business structure, they are likely to be noncompetitive. Without appropriate integration, the key advantages of electronic commerce risk becoming fractionated, lost, and unaccounted for, as they were in the isolated technical republics in the information archipelagos of the past.

A basic strategic information technology vehicle in the age of electronic commerce is the emergence of an enterprise infrastructure that integrates previously disparate systems and behaves as an organic information nucleus for the enterprise. For businesses, the unconnected industrial information hierarchy of the past will be deconstructed in the electronic commerce workplace evolving to a new interconnected, information-based organizational form. Just as a nervous system allows an organism to behave in a coordinated fashion that optimizes its senses, strength, mobility, and defenses for survival, the new electronic commerce infrastructure will provide the data and information nucleus that will help collectively guide, protect, and grow the organization of the future.

The open architecture of the Internet and intra-company intranets and external extranets are playing lead roles in stimulating this critical evolutionary achievement.

3. Enterprise Extension

The third natural progression in the electronic commerce evolutionary process is the electronic extension of the enterprise and the formation of the extended business community. Community formation occurs among groups of closely dependent companies and, ultimately, consumers will drive powerful changes in the enterprise and inter-enterprise business rules for existing companies. With the deployment of an electronic commerce internetworking environment, traditional organizational boundaries begin to blur, and businesses begin to integrate processes and data with customers, suppliers, and other business partners. The inevitable transformation that such integration will drive replaces the individual linkages on the supply chain of "supplier," "manufacturer," "distributor," "retailer," and "consumer" with a more tightly integrated link of "community," where individual members rely on mutual cooperation for success.

In the coming competitive environment, extended electronic communities will enable leading companies to maintain very low levels of inventory and execute rapid product cycles, where products can be built, configured, and replenished to order. These pioneering companies will integrate across the entire supply chain and work together as a virtual electronic commerce team. For top performing communities, the bonding of manufacturer, distributor, retailer, and consumer will occur via a common electronic commerce infrastructure; not connecting to the community will be a catastrophic competitive loss.

The market forces driving electronic commerce are so compelling because businesses will not be able to compete without it. As community models are adopted, the new electronic infrastructure will allow companies to manage their entire operational flow across the supply chain. In doing so, businesses will enhance their customer relationships, achieve leverage in the marketplace, reduce costs and the need to maintain large inventories, and obtain faster and more reliable deliveries of materials and services. The strategic value of dynamically linking the various processes of requisitioning, buying, transporting, receiving, paying for, and managing material services is that it transforms traditional relationships from adversarial to cooperative. In the age of electronic commerce, there is truly shared risk and reward.

The Last Mile: Where Telecommunications Traffic Slows to a Crawl

Don M. Flournoy, Ph.D.
Professor
Institute for Telecommunications Studies,
Ohio State University

Thomas N. Scott
Principal
The Veda Home Company

Abstract

This paper addresses the problem of congestion on access ramps of the information superhighway and examines both near- and longer-term solutions. The focus is on customers and their ability to gain access to multiple transit networks and service providers by building bridges into and out of the customer premises and by managing information within the home. A new type of service/connectivity provider, a local services integrator (LSI) operating in the customer loop, and a customer premises integrator (CPI) operating within the residential environment is recommended. An LSI may be either the operator of a multiple services network such as an LMDS licensee; a digital-ready CATV operator running over a hybrid fiber/coax infrastructure; a phone company or CLEC using multimegabit DSL technologies; or a heretofore unknown entity with the capital, technology, and regulatory authority to negotiate services from any and all providers and to integrate them into an affordable, easy-to-use residential package. The mission of the CPI is to bring to market innovations for the managing of telecommunications traffic within the home. Since the mix of services will be both symmetric and asymmetric—arriving from more than one service provider using multiple signal formats deployed over a growing number of wireline and wireless media—home applications must be dynamically configurable and controlled by both CPI and customers.

The Problem

Telecommunication highways of terabit-per-second speeds now link most American cities but few homeowners have access to even a megabit of speed. Fiber stops before it gets to the home. Coaxial cable TV (CATV) provides broadband service but, with few exceptions, is still a one-way pipe. Telephone lines are slow and inefficient.

In the wireless world, local stations soon will begin digital TV (DTV) broadcasting in competition with direct-to-home broadcast satellites (DBS), broadband wireless such as LMDS and MVDS, and other providers aimed at luring residential consumers as subscribers. Yet, in this complex offering of digital multicast programming, only very modest prospects for interactivity are now possible, and integration of the various wireless and wireline services is still some distance away.

From the consumer's perspective, the problem does not seem to be a lack of future options. On the contrary, the consumer is promised more choice, more bandwidth, and more speed than previously imaginable. But which one—which ones—to choose? Will the ISDN-anywhere lines planned to connect to corporate networks also enable access to the Internet? Will the cable TV, satellite TV, and over-the-air broadcast channels that feed the home entertainment center offer more than passive viewing? Will any of these technologies be integrated with each other? In brief, will the enabling technologies be in place and will the services be packaged in a way that the family can both afford them and get maximum use out of them?

The number of suitors is certainly growing, as are the number of telecommunications paths into the home and the variety of services planned. Awaiting the go-ahead from those who make the business decisions are a surprising number of technological solutions capable of addressing the traffic jams between the end of the nearest high-capacity trunk line and the home premises. Breakthrough developments in hardware and software in the wireline and wireless domains have given telecommunications managers new tools with which to work.

What seems to be focusing attention on the home consumer is the prospect of selling increased-speed access to the Internet, switched video services such as video-on-demand, and, to an extent that is still undefined, easy exchange of multimedia products among end users. Whether these are niche markets—therefore limited—or a global mass market waiting to take off remains to be seen. Technological developments will represent a big part of whatever solutions there are, but the answers will also be largely economic and part regulatory, social, and cultural. As long as the last-mile infrastructure is absent or unresponsive to the felt needs of those consumers, however, no serious market will develop.

A Proliferation of Paths and Players

All customers have their own reasons for reaching into that vast store of information and interactivity that has come to exist in the digital world and bringing it into their home. These could be reasons related to communicating with others, connecting to work, engaging in learning activities, or just enjoying life. In brief, consumers are open to sales pitches from service providers.

Soon there will be multiple broadband corridors heading in the direction of the home, a fact which consumers will not realize by talking to only one vendor. Providers of information services are already engaged in research on the best ways to use one or more of these paths to help consumers realize their dreams. What the vendors will not tell customers is that for the foreseeable future both they and their service providers will be wringing their hands over the inadequacies of a telecommunications infrastructure that promises much but delivers little. These channels, which will be wireless as well as wireline, will all be incompatible. Not only will they require specialized interfaces to the terminal equipment within the home, they will require different access technologies beyond the home.

Wireless and wireline access providers have traditionally operated in different worlds. Historic territorialities and regulatory oversight have dampened efforts at collaboration. Each has represented protected industries with its own consumer base that was huge and profitable. Their subscribers / audiences have now fragmented and their missions are less well-defined; competitors have emerged.

Another reason for this complexity is that delivery of interactive broadband services to the home is a recent phenomenon largely occasioned by the Internet. Until now, research and marketing departments have worked in vain to come up with a clear picture of what the next generation of videotext, videophone, movie-on-demand interactive data services would look like, or for what consumers would be willing to pay. Because profitable services are already in place, there is massive anxiety over leaving the safe havens of voice (telcos) and one-way video (broadcasters) and fear of encountering some unforeseen disaster in the cutover from analog to digital. One way to think about today's customers, providers, and delivery / access networks is to use the model developed by the International Telecommunications Union (see *Figure 1*).

The model helps analyze and simplify the complexity of services, players, and paths into the home. This model can be applied to any end-to-end communication regardless of the underlying medium: wireless or wireline, narrowband or broadband.

For example, in a plain old telephone service (POTS) call:

• *A* and *B* are two customers talking on the phone.

• The customer-access points-of-interface (POI) are the standard network interface (SNI) boxes attached to the homes.

FIGURE 1

ITU Reference Model

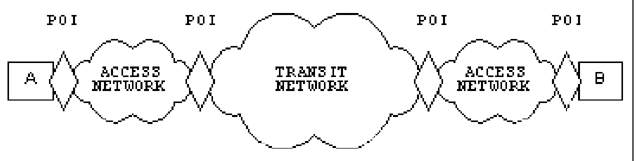

POI: POINT OF INTERFACE

• The access networks are the copper loops connecting the SNIs to the central offices.

• The access / transit POIs are the central offices. In certain analyses, the POI could be equipment in the CO such as the main distributing frame or switch.

• The transit network is the public switched telephone network (PSTN) of long-distance, interexchange carriers.

Wireline Access

The following wireline technologies provide the home and small business with digital access to the Internet and other interactive services:

• SW56 / ISDN, switched 56 and integrated services digital network

• HDSL / SDSL, high-bit-rate DSL and symmetric / single-line DSL

• ADSL / VDSL, asymmetric DSL and very-high-bit-rate DSL

• FITL, fiber-in-the-loop

• HFC, hybrid fiber / coax

SW56 and ISDN

The first access networks were those of the telephone companies, consisting of analog signals carried entirely on twisted pairs or loops of copper. Later refinements of the customer loop used analog and digital pair-gain technologies to multiplex lines coming out of the central office, but the final drop into the home remained the analog copper pair. Analog telephony was then augmented by digital signals into the home, ushering in the digital subscriber line (see *Table 1*). Early versions of DSL technology were switched 56 (SW56) and integrated services digital network (ISDN).

Residential ISDN is not available everywhere, and it is expensive to install and subscribe to, but such lines can provide twice the speed of telephone lines using the fastest (56 kbps)

TABLE 1

DSL Abbreviations

Abbreviation	Full Name
DSL	digital subscriber line
xDSL	x = whatever DSL
SW56	Switched 56
ISDN	Integrated Services Digital Network
IDSL	ISDN DSL
HDSL	high-bit-rate DSL
HDSL2	HDSL version 2
SDSL	symmetric single-line DSL
ADSL	asymmetric DSL
RADSL	rate-adaptive asymmetric DSL
VDSL	very-high-bit-rate DSL

analog modems available today. With the discovery of faster DSL technologies in the late 1980s by a research team led by Joseph Lechleider at Bell Communications Research, phone companies had the ability to convert their copper outside plant into multimegabit access networks.

HDSL and SDSL

After SW56 and ISDN, the next DSL technology to be deployed by the phone companies was high-bit-rate DSL, which runs over two twisted pairs of copper. The main application of HDSL has been to replace T1 lines, the high-capacity (1.544 Mbps) workhorse of local trunking. Because the T1 carrier requires repeaters at intervals of 3,000–6,000 feet and tuning of these repeaters is labor-intensive and therefore expensive, repeaterless HDSL was embraced as soon as it was ready for deployment in the early '90s. Since 1996, a one-pair version of HDSL, called HDSL2 or symmetric single-line DSL (SDLS), has also been available.

An instructive application of HDSL and SDSL is in the cellular telephone network (see *Figure 2*). In this example, an HDSL line connects the cell site to the mobile telephone switching office (MTSO), which in turn is connected to the nearest telco central office by either another HDSL / SDSL line or by fiber (if PSTN traffic between the MTSO and CO warrants the investment).

ADSL and VDSL

In outside plants where copper extends from the central office to the customer premises, the asymmetric digital subscriber line (ADSL) is an obvious choice for the telco or CLEC wishing to enter the broadband access market. A likely scenario that will reach a significant number of residences is shown in *Figure 3*.

In this example, a network access provider (NAP) installs a wire center, consisting of ADSL modems or ADSL transmission units-central office type (ATU-C) and a DSL access multiplexer (DSLAM). The NAP connects the wire center to various transit points-of-presence in or near the central office; POPs are the interface between the residential customer's access and transit networks. Network service providers such as ISPs and switched video service providers connect to transit through ac-

FIGURE 2

Technology Applied to Cellular Infrastructure ADSL and VDSL

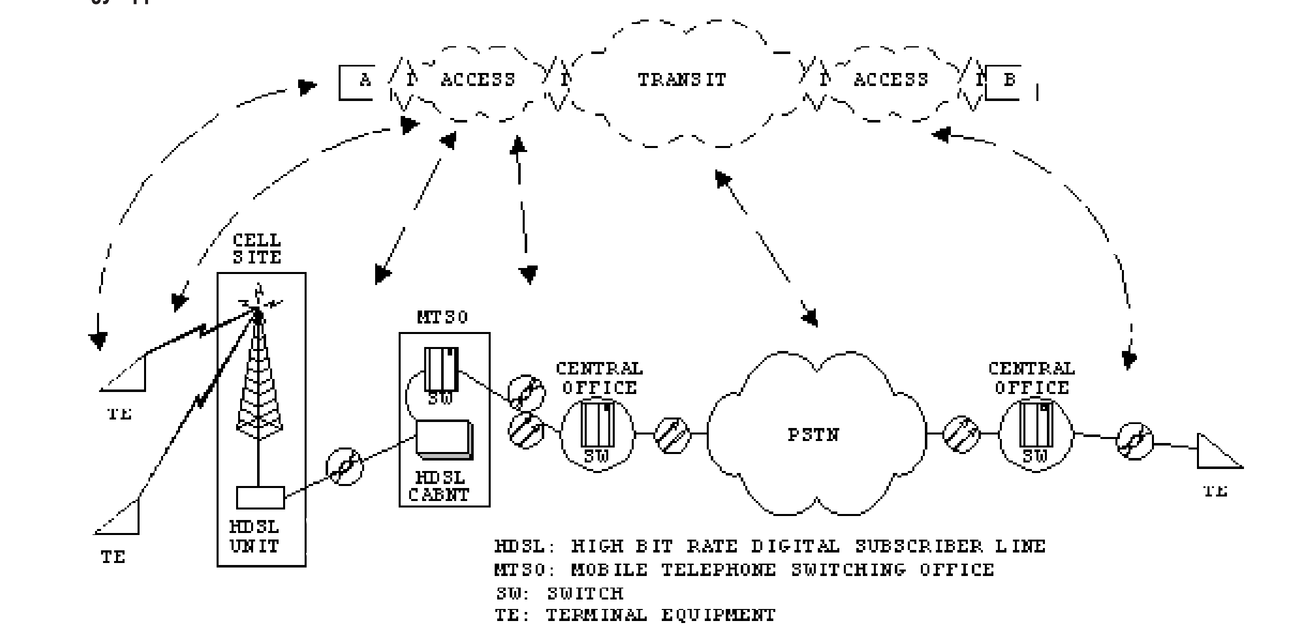

FIGURE 3

ADSL Reference Model

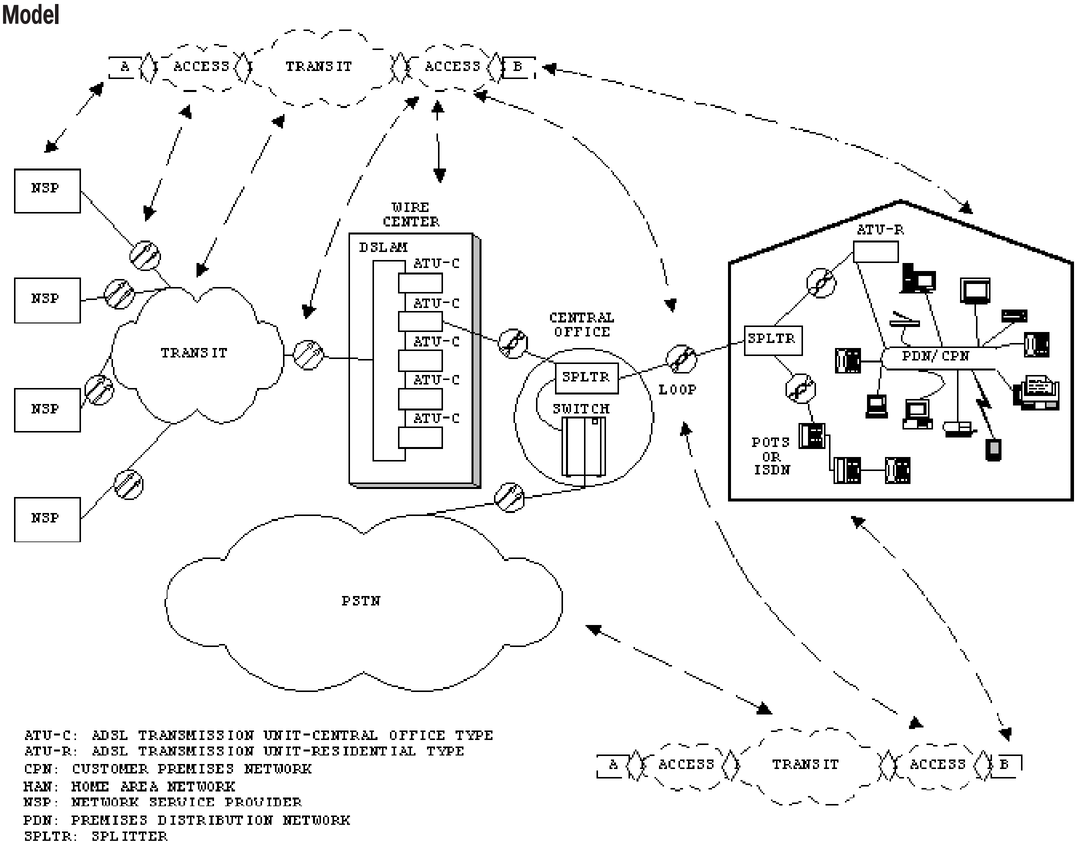

cess networks on their side, thus completing connectivity between customer *A* and customer *B* in the ITU reference model.

The wire center may be colocated in the central office or in a facility nearby. A customer premise terminal within the service area of 18,000 feet from the wire center communicates with the ATU-C over the copper loop by means of a residential ADSL modem or ADSL transmission unit-residential type (ATU-R). Downstream transmission rates are in the 2–8 Mbps range, upstream 64-1,000 kbps.

Very-high-bit-rate DSL (VDSL) is a high-speed, short-range version of ADSL. For the moment, the choice between ADSL and VDSL is a trade-off between distance and transmission rate. The longer the distance, the slower the transmission; in such a situation ADSL technology that can be upgraded easily and inexpensively to VDSL would be a good choice. In a hybrid ADSL/VDSL system—a broadband extension of rate-adaptive DSL (RADSL)—currently under development, projected rates are 52 Mbps in the VDSL range over shorter distances and 2–8 Mbps in the ADSL range over longer distances.

FIGURE 4

POTS, ISDN, and ADSL Frequency Spectrum Fiber-in-the-Loop

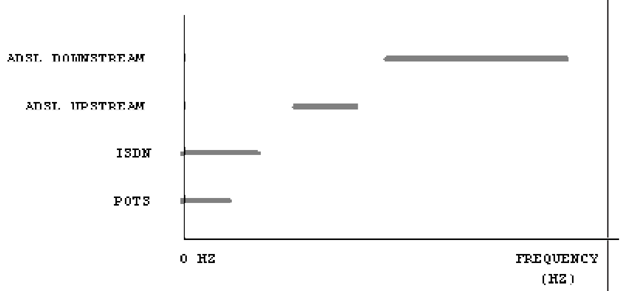

Comparison of SW56/ISDN and ADSL/VDSL

Switched 56 and ISDN pale in comparison to ADSL and VDSL transmission rates. However, the former has proven track records and effectively serves the needs of many residential applications. As researchers and engineers amass experience with the newer, faster DSL technologies and as competitive LECs enter the market as NAPs and set up wire centers near each CO, it is reasonable to expect ADSL or VDSL to be deployed virtually everywhere in the customer loop.

A valuable feature of ADSL is its compatibility with POTS. Splitters shunt the lower-frequency POTS signal to the PSTN on the network side of the loop and to plain old telephones on the residential side (see *Figure 4*). Since POTS power is fed from the CO, the ADSL customers always have dial tone for 911/lifeline service if the customer maintains a legacy phone. This is not the case with off-the-shelf ISDN for two reasons. First, POTS and ISDN cannot use the same twisted pair. Second, ISDN customer premises equipment must have their own power feed. If customers lose power, the BRI line goes dead—unless of course customers have battery or generator backup.

Fiber-in-the-Loop

Fiber-in-the-loop (FITL) comes in several different varieties:

- Fiber-to-the-home (FTTH)
- Fiber-to-the-curb (FTTC)
- Fiber-to-the-neighborhood (FTTN)

Because of its bandwidth and robust signal, fiber-to-the-home (FTTH) is an ideal medium (see *Figure 5*). But FTTH carries such a high price tag that no one gives it serious consideration in the near term. Less expensive, more practical access networks can be built by extending fiber from the central office to an intermediate point, where an optical network unit (ONU) converts the signal from optical to electrical, which is relayed on existing copper lines into the home. The section of the network that connects the ONU and central office is the digital loop carrier (DLC). When the ONU serves a large number of homes, the access topology is fiber-to-the-neighborhood (FTTN); when the ONU serves a smaller number of homes, it is fiber-to-the-curb (FTTC).

When the distance from the ONU to the customer premises is relatively short, as is the case with FTTC, VDSL is a feasible option. Longer distances in a FTTN topology can be served by ADSL but the result is slower transmission. Long and short are relative terms, however, in today's fast-moving technological environment, as are channel capacities and rates of transmission. Hence, a distance that supports ADSL transmissions under current technological constraints may soon be suitable for VDSL.

CATV and HFC

Just as hybrid fiber/coax (HFC) will help telephony providers upgrade their residential services to include broadband communications such as data and video, HFC will help CATV operators add bandwidth and interactive services to their portfolio (see *Figure 6*). This new cable architecture, in which information arriving in the neighborhood on fiber is carried into the home on coaxial cables, will support multi-megabit Internet, VOD, games, shopping, and other interactive services as well as telephony. In addition, the infrastructure can be upgraded to support broadband telephony services such as videophone and videoconferencing as well as digital HDTV.

The key to conversion of cable television into an advanced digital network is the installation of new high-speed modems in the subscriber premises and modifications in the cable plant that permit digital signal compression, encryption coding/decoding, and real-time two-way data communications based on Internet protocols. Although cable modems promise to deliver downstream speeds 100 times faster than ISDN modems, rollouts of these systems have been slowed by lack of compatible standards and unrealistically high installation costs.

Comparison of DSL and HFC

A communication system is no faster than its slowest link. Although the capacity of the fiber trunk of an HFC system is for all practical purposes unlimited, the last leg or feeder and the drop that enters each home is metallic (coax) and therefore limited. The coax feeder shares its bandwidth with all customers connected to it. One hundred customers who are Web surfing and pulling in separate VOD streams can degrade quality-of-service parameters to an unacceptable level even if feeder speeds of 10–20 Mbps or higher are attained.

This is not the case with ADSL and VDSL. All customers have their own dedicated line that is fed from an ONU or directly from the CO. The network-side fiber feeds to the ONU and CO can be configured to support any customer demands, provided that the access operator can afford the switching and multiplexing electronics and the service provider's transit network is adequate to the task.

Wireless Access

High-speed access to Internet and other broadband interactive services may be provided to the home via, or in conjunction with, wireless telecommunications channels in some of the following formats:

- MMDS, multipoint multichannel distribution system
- LMDS/MVDS, local multipoint and microwave video distribution system
- WLL/RITL, wireless local loop/radio in the loop
- IVDS/ITV, applications of interactive video and data services in interactive TV
- VSAT/GEO/LEO, very small aperture terminals communicating with GEO/LEO satellites

MMDS

Providers of wireless cable, the omnidirectional multipoint multichannel distribution system (MMDS) using microwave (2.5–2.7 GHz) technologies with a compressed digital signal, are moving rapidly to install return channel capability to their formerly one-way broadcasts.

The principal advantage of the wireless cable approach as an element of the access loop is the availability of big chunks of under-utilized spectrum. These segments of spectrum will be more valuable and flexible once digital kicks in. MMDS services have been established in a decade or two. There is a wealth of experience with this previously one-way distribution technology. System implementation, which is little more than an installed transmitter on a high tower and a small receiving antenna on the customer's balcony or roof, is quick and inexpensive. Although MMDS has been slow to find its place in U.S. markets crowded by broadcast TV, DBS, and cable, the broadband technology is popular in cities such as Beijing, Mexico City, Nairobi, Riga, and Moscow.

With the deregulation of U.S. telcos, Bell Atlantic, NYNEX, Pacific Telesis, and others invested heavily in MMDS franchises as a way to get up and running as full-service providers. By 1997, a major telco-programming venture (Tele-TV) had failed and several of its underwriters slowed the buildup of MMDS until a more suitable strategy for approaching the broadband home market could be found.

FIGURE 5

Fiber-in-the-Loop

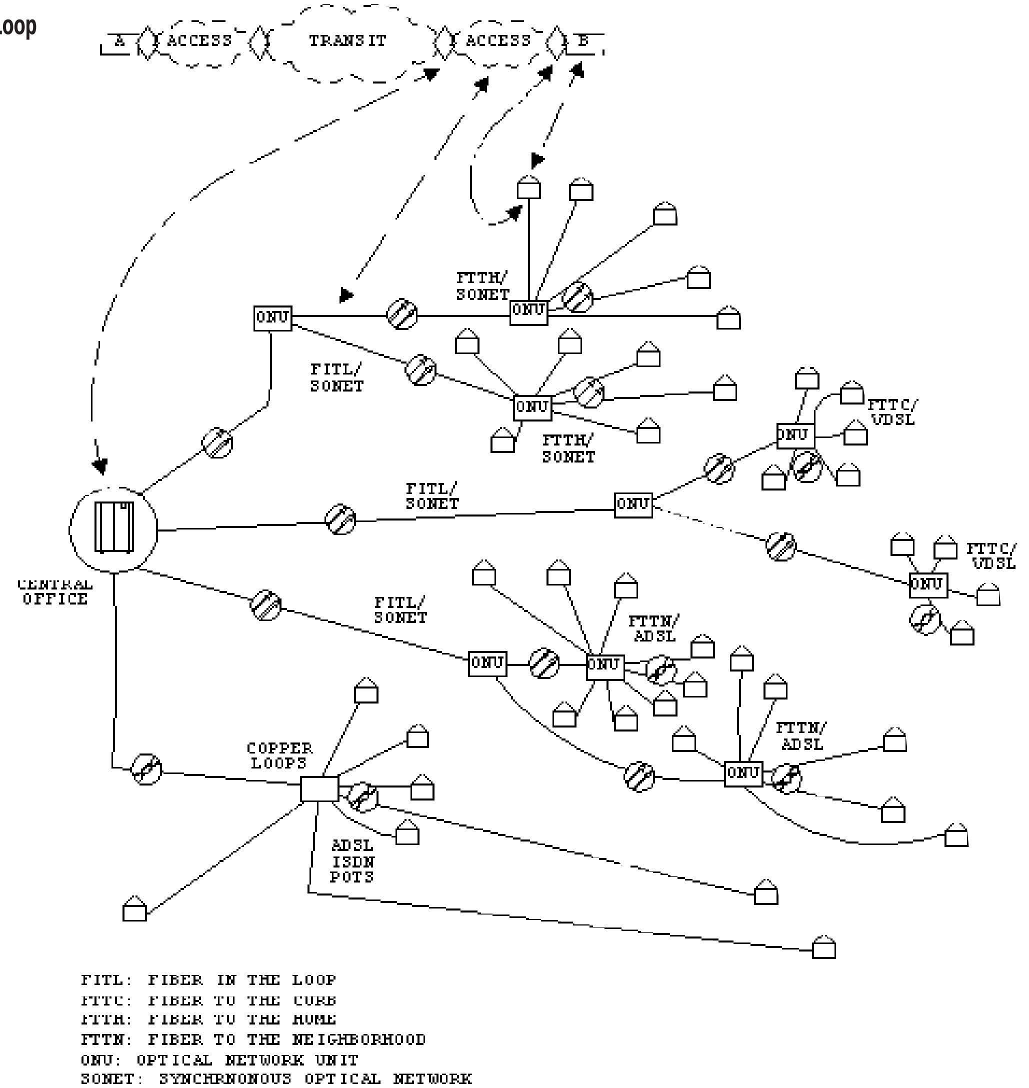

FIGURE 6

Hybrid Fiber/Coax Architecture

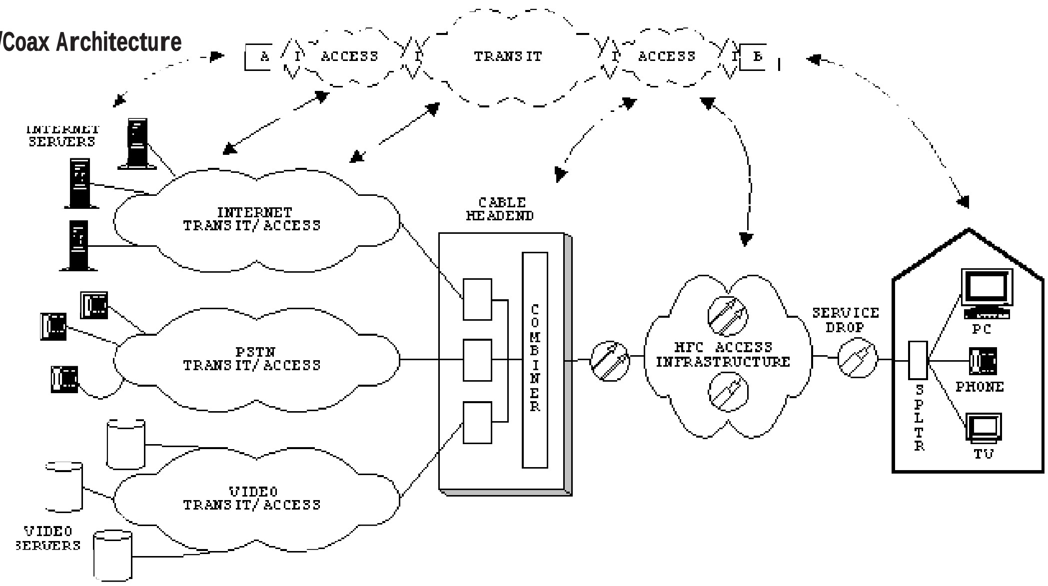

FIGURE 7

Local Multichannel Distribution System

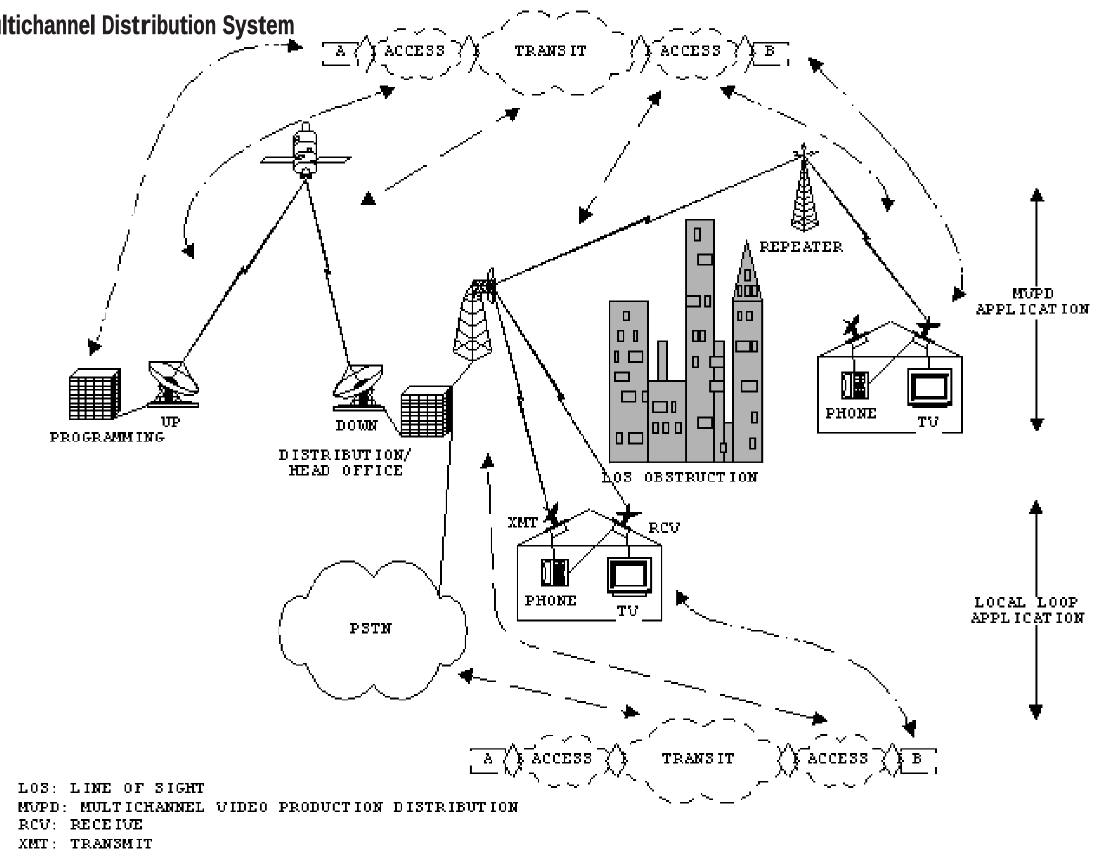

These actions caused a loss of investor confidence, MMDS stocks were hard hit, and some momentum in developing this technology has been lost.

LMDS and MVDS

Colloquially called cellular TV, the 1997-approved local multipoint distribution system (LMDS), operating in the upper microwave (28–31 GHz) frequencies, is designed to be an interactive broadband service using interconnected cells in the local loop (see *Figure 7*). Each transmitter serves an area of 2–3 km in diameter. Cell architecture is similar to that of cellular radio and personal communications services, but LMDS cell size is significantly smaller.

In auctioning the LMDS spectrum, the FCC has designated two licenses in each of the 492 basic trading areas (BTA) in the United States. The commission allocated a huge block of spectrum—1150 MHz for Block A licenses, 150 MHz for Block B—sufficient to provide telephone, video, and data services. With its two-way transceiver capabilities, LMDS is a promising medium for Internet access, videophone, videoconferencing, and pay-per-view (PPV) cable television.

LMDS is sometimes referred to as "wireless fiber"to emphasize its greater bandwidth in comparison to wireless cable (MMDS). Engineers estimate that LMDS technologies can produce 1 bit per second for each Hz. Hence, a Block A spectrum of 1150 MHz is capable of transmission rates in the 1 GHz range, more than enough bandwidth for simultaneous support of the broadband services discussed in this paper.

A limitation of LMDS is the modest experience base for the technology. Early indications are encouraging. The FCC granted a pioneer preference license to CellularVision of New York. According to reviews, the digital service performs admirably. However, LMDS has not yet been tried in many locations in the United States, so whether the long-term performance of the technology rivals that of MMDS and PCS remains to be seen.

It is conceivable for a Block A operator to offer narrowband voice, video, and Internet access as well as broadband videophone, videoconferencing, multimegabit Internet, and interactive multichannel digital HDTV—and do it on one unified infrastructure of hubs and subs. Realization of this scenario is less likely to hinge on the technology than on economic, marketing, and regulatory factors.

The commission left the door open to spectrum disaggregation and geographic partitioning. With few limitations, a licensee can sell smaller bands of its spectrum while partitioning the BTA into geographical units to be sold individually. Will each BTA witness a well-funded, unified LMDS slugging it out with incumbent telcos and cable, or will the spectrum and geography be balkanized into bits of minuscule but profitable services? Will major metropolitan areas be more supportive environments for unified service than their small town and rural counterparts? Answers are forthcoming in the marketplace.

The European version of LMDS, microwave video distribution system (MVDS), transmits at frequencies above 40 GHz and has a signal range similar to that of LMDS. Although the higher frequencies at which MVDS operates are less crowded and therefore capable of supporting more services, not only are the engineering challenges more difficult to meet but the manufacturing and maintenance of such systems pushes investors and network operators to the extreme. However, today's problem at 40 GHz is tomorrow's piece of cake. Technology marches on.

WLL/RITL

The wireless local loop (WLL), also called radio-in-the-loop (RITL), adds its own unique set of contributions to the improvement of telecommunication capabilities both outside and inside the customer premises (see *Figure 8*).

WLL applications are most commonly found in developing countries where wireless approaches provide a rapid and cost-effective way of delivering telephone services to not-yet-connected homes and businesses. Communications services are often made available in remote communities by connecting them to the nearest central telephone office via microwave carrier, satellite, or other means. At the local point of presence, radio ports are installed that are interfaced to users' home or business equipment.

WLL is currently being tested as a bypass solution and alternative to existing wireline networks in more advanced societies, including tests of broadband transmissions capable of supporting multimedia exchange. Among the applications are switched telecommunications services for small and medium-sized business customers providing a single source for local and long-distance telephone and Internet and other high-speed data access as part of the nationwide deployment of competitive local exchange (CLEC) services in the United States. These wireless carriers are giving home and business customers a single point-of-contact for multiple services. One such company uses previously unavailable 38-GHz frequencies to provide its local communications services.

Applications of IVDS in ITV

The interactive video and data service (IVDS), a new return path for making broadcast TV interactive, was approved for FCC spectrum auction in 1994. IVDS receiver/transmitters are installed in subscriber homes in the form of set-top boxes (see *Figure 9*). Each box hosts a wireless radio frequency modem with remote control through which customers can interact with the programming source. Strategically located cell-tower transceivers gather signals from interactive video users and relay them via satellite or terrestrial links to the transmitting station or affiliated office.

This technology is being positioned as a way to give added value to the TV broadcasts of local stations. With the prospect of a viable home-to-station return path in the local market, IVDS will be an option under consideration as U.S. television stations look for ways to profit from their new DTV channel assignments.

One of the limitations of this option is that the frequencies were auctioned before the technology to use them had been perfected. Set-top boxes were late in arriving, and the FCC found itself facing multiple defaults by auction winners. Such defaults do not diminish the long-term value of the technology, but they have made investors cautious.

VSAT Applications in GEO/LEO Satellite Systems

This is an interactive satellite service (see *Figure 10*). It gets its name from the very small aperture terminal (VSAT) dishes found on business premises, now being seen on home premises as well. The VSAT service operates from geosynchronous orbit (GEO) relaying digital data, voice and video signals, in real time or downloaded for later use, to on-site computers. The upstream is by return signal to the satellite or, as in the case of the Hughes DirecPC home and small business service, Internet searches and requests are made by way of low-speed terrestrial link (telephone lines) with content delivery direct to user premises via the advanced digital capabilities of satellite.

Rapid implementation of service is one of the attractive features of VSAT technologies. Installation of subscriber equipment—the dish, the receiver, and computer interface—can be accomplished in a day or two. With mass marketing and ready availability of digital DBS services, costs are coming within reach of users who have a need to download Internet data at faster speeds. The technology offers an advantage over cable modems and digitized phone lines in that it is available everywhere, whether in rural or urban areas.

There are constraints in the hybrid satellite/twisted pair telephone line configuration. The customer can order up material from heavily loaded Web sites and great quantities of text, graphic, audio, and video information can be delivered quickly. However, if users wanted to send large files, the outpath is likely to be the same slow telephone lines that forced them to the satellite option in the first place. Direct return to the satellite of home-originated transmissions is possible but for the moment not yet a product for mass consumer use.

More than one low-earth orbiting (LEO) satellite system is now approaching launch. These constellations of satellites that communicate with terrestrial stations and among themselves will orbit some 800 kilometers or higher and provide on-demand global Internet access, videophone, videoconferencing, and interactive multimedia to fixed and mobile transceivers. LEO services are suited for urban or rural areas not connected yet to a broadband terrestrial infrastructure or one that cannot be covered economically using traditional terrestrial infrastructures. These year-2000 plus technologies are suited especially to bidirectional asymmetric services as they offer a very short round-trip propagation time, enabling them more easily to share common communication protocols, standards, and applications with terrestrial networks.

FIGURE 8

Wireless Local Loop/Radio in the Loop

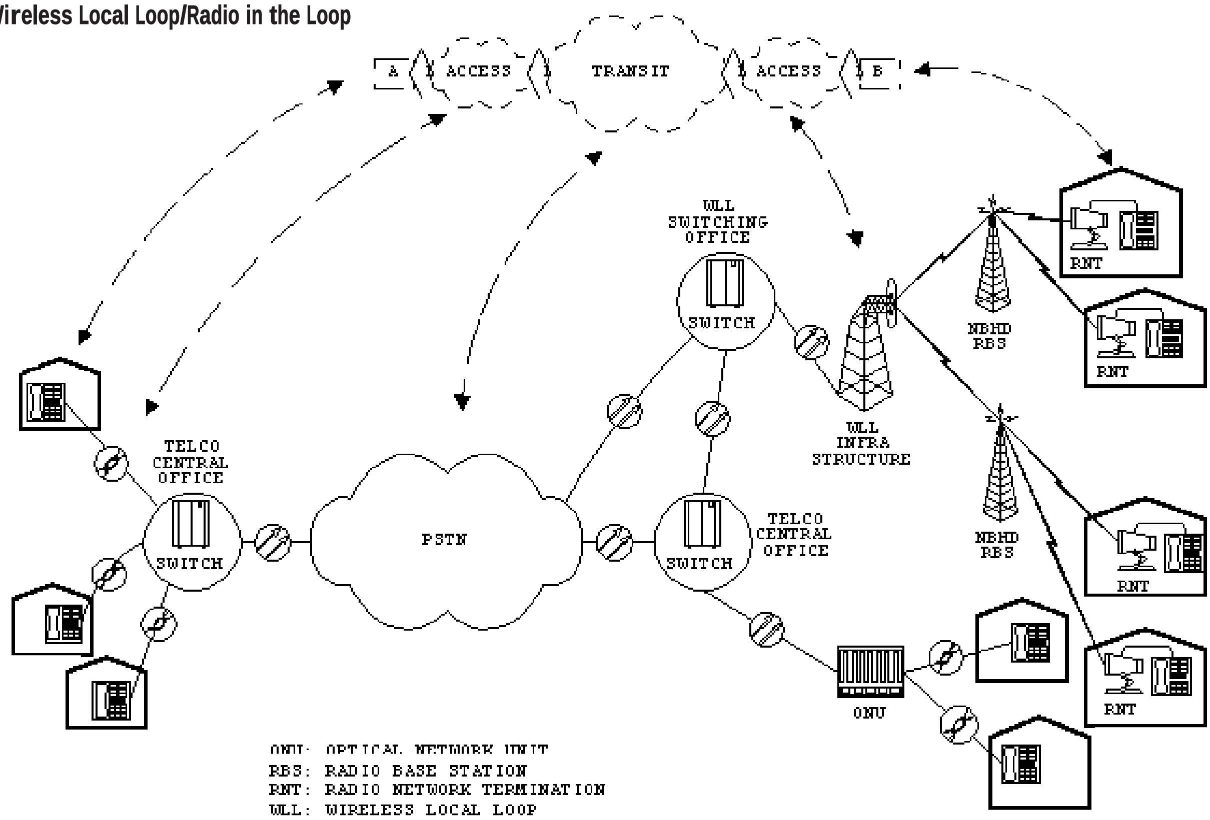

FIGURE 9

IVDS Interactive TV

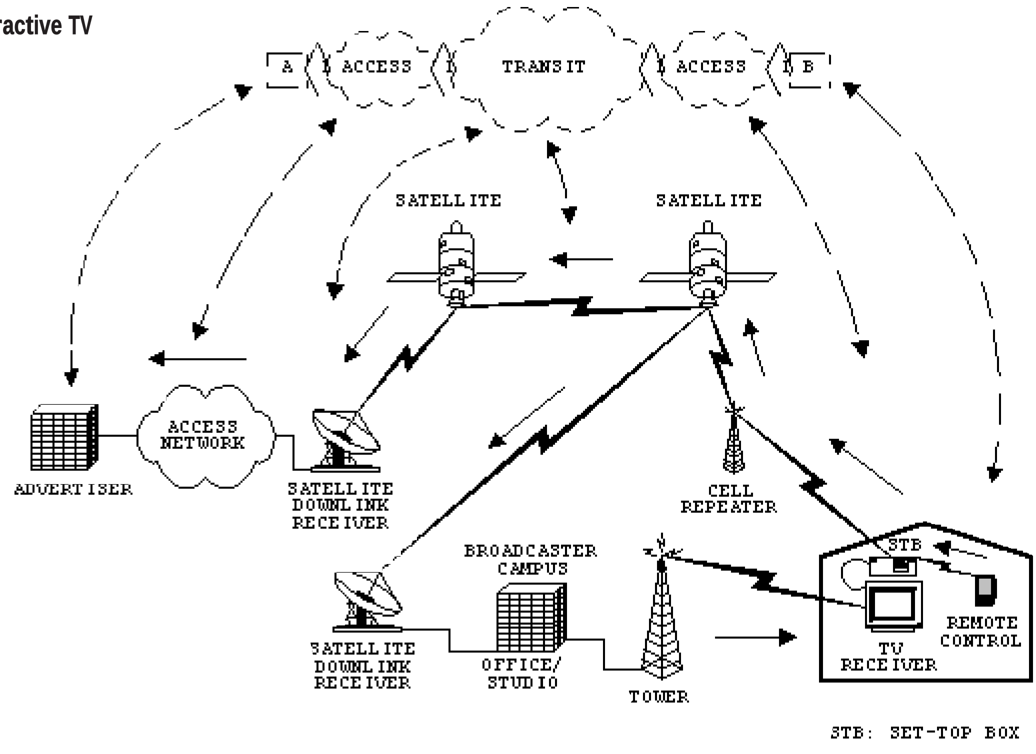

FIGURE 10

VSAT Applications in GEO and LEO Satellites

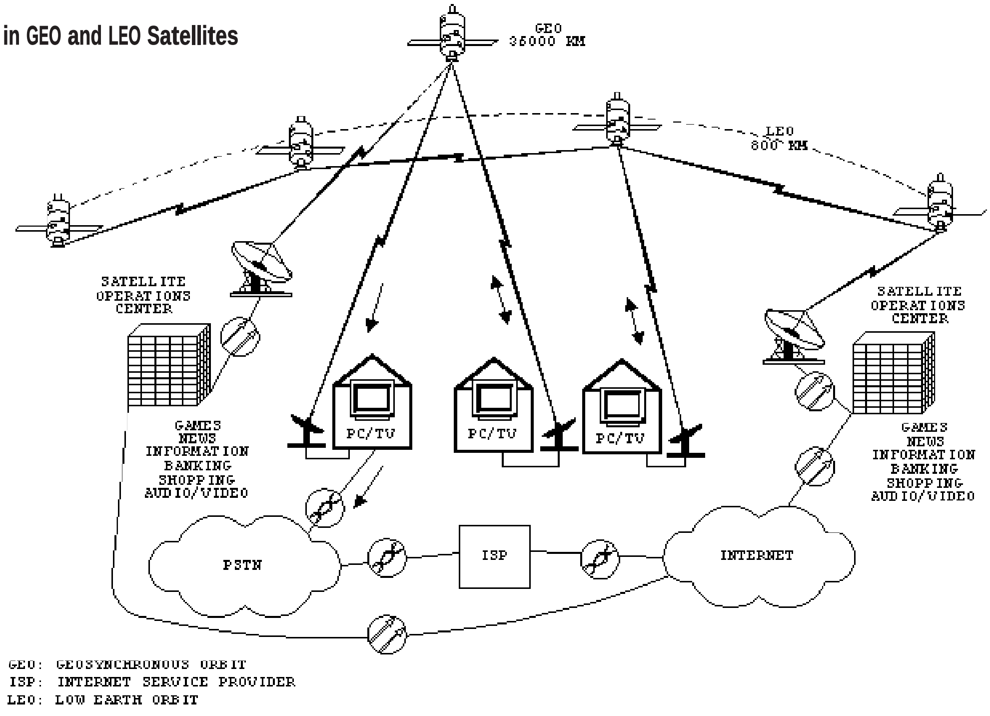

Convergence in the Local Loop

A quick R&D literature review and visits to various telecommunications forums will show that we truly are on the brink of plunging into broadband residential services running on multiple networks comprised of components from multiple manufacturers and vendors. The reason for these industry initiatives is the realization that there is a need to accommodate those services defined by third-party integrators who do not operate the networks over which the services run—and that a changed regulatory environment will permit this to happen.

Providers of home services and those working on near-home connectivity have been held back by the fact that service architectures have been dependent on a particular core transport network, either on a particular transport technique or bearer service. Due to the multiplicity of services, technologies, and standards, it is difficult to port services (see *Figure 11*).

What is needed is a more open software creation platform so that multiple services from multiple providers have a chance of being offered transparently across multiple networks, which would enable individual services to be constructed and rapidly deployed using systems and software developed by different vendors. Thus, the retailer of services, the Internet services provider, and the games-on-demand operator are all able to integrate their systems into a plug-in environment, making more feasible seamless communication of feature-rich information into and out of the customer premises.

Supported both by technical and regulatory changes, we are likely to see the impossible happen now. Bearer services will begin to cooperate, allowing their diverse networks to interface and be configured by independent service / connectivity providers in the last mile. This will happen for a simple reason: nobody wins when customers are confused, frustrated, and stretched beyond their ability to pay. Given the following, standardization—the only form of cooperation likely to succeed in the aggressive, competitive telecommunications industry—is the best way to grow the business:

- Telco, CATV, MMDS, DBS, and now LMDS and DTV operators are all chasing the same residential customer.
- Customers would like to have a small part of all of them.
- Consumers have the monthly resources to cough up full subscription fees only to one or two.

How might this cooperative competition come about? One prospect is that a single powerful operator with enormous capital and investor backing, persuasive technology, and a bundle of services will capture and dominate the last mile market, and this could happen. A different model is both more desirable and more likely. The model is that of the local services integrator (LSI), who will provide robust, convenient, and affordable connectivity to a myriad of services found on diverse networks. Which technology or technologies will be used? There are several promising possibilities including ILEC / CLEC using xDSL, CATV using hybrid fiber / coax, and LMDS / WLL using wireless fiber.

Customer Premises Integration

There is also need for integrative services within the home (see *Figure 12*). Bellcore and others have done extensive research on home area networks (HAN) and cus-

FIGURE 11

Telecommunications Options in the Local Loop

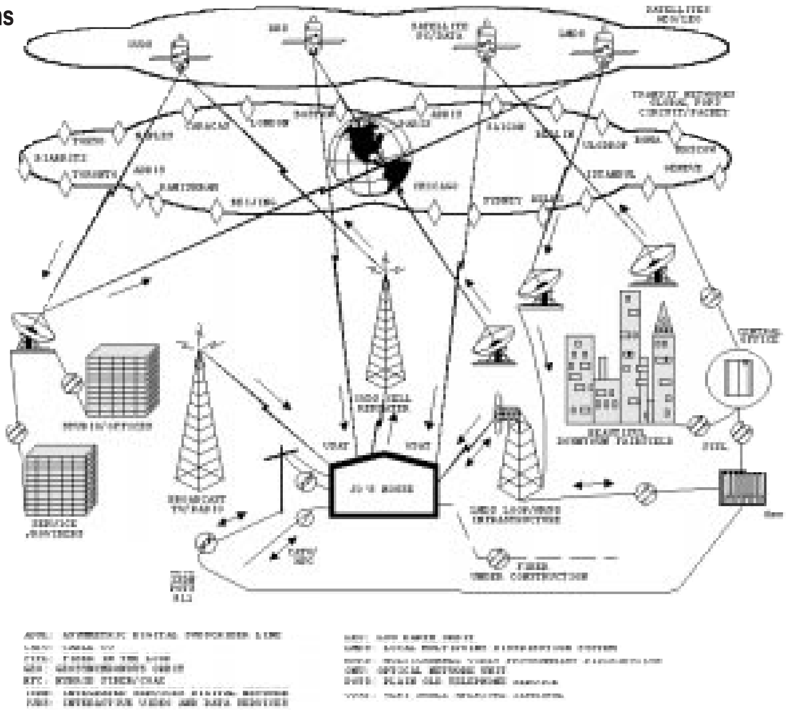

FIGURE 12

Teleommunications Options in the Customer Premises

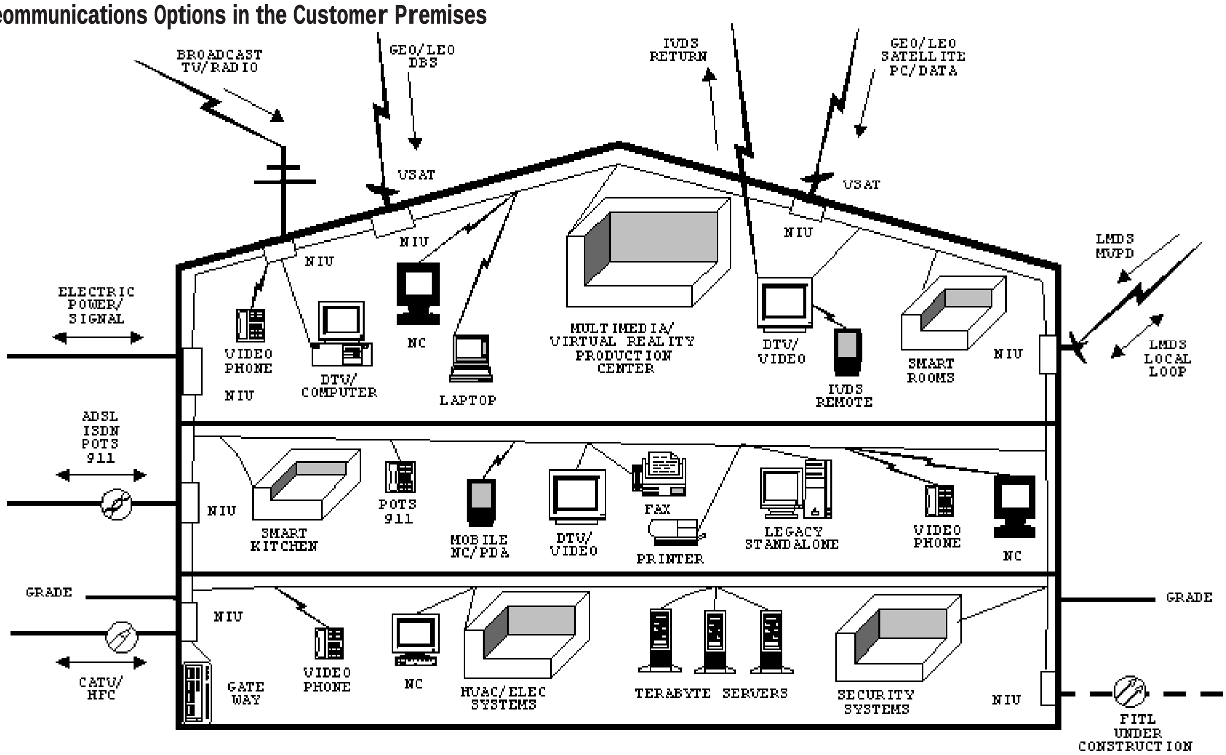

ADSL: ASYMMETRIC DIGITAL SUBSCRIBER LINE
CATV: CABLE TV
DTV: DIGITAL TV
FITL: FIBER IN THE LOOP
GEO: GEOSYNCHRONOUS ORBIT
HFC: HYBRID FIBER/COAX
HVAC: HEATING VENTILATION AIR CONDITIONING
ISDN: INTEGRATED SERVICES DIGITAL NETWORK
IVDS: INTERACTIVE VIDEO AND DATA SERVICES
LEO: LOW EARTH ORBIT
LMDS: LOCAL MULTIPOINT DISTRIBUTION SYSTEM
MVPD: MULTICHANNEL VIDEO PROGRAMMING DISTRIBUTION
NC: NETWORK COMPUTER
NIU: NETWORK INTERFACE UNIT
PDA: PERSONAL DIGITAL ASSISTANT
POTS: PLAIN OLD TELEPHONE SERVICE
VSAT: VERY SMALL APERTURE TERMINAL

tomer premises networks (CPN). The goal has been to work toward common standards for interconnecting, adding to, and upgrading the services of multiple providers in the broadband home market, thereby simplifying access and giving users greater control over traffic within the home.

Here, too, as with the local services integrators, the challenge is to devise a means by which home users may gain on-demand access to the networks of more than one operator. But what is even more demanding of this model is the need to give enhanced creation, storage, distribution, and control capabilities to customers and their families so they can customize applications to taste. The customer is in the driver's seat. This assumes that each member of the family, according to that person's recreational, educational, professional, or other agenda, will be free to roam from one network to another, connecting, sampling, downloading, clipping, pasting, processing, storing, exchanging, and consulting at will. Ideally, this approach will account for service requests that could originate at any time from any home terminal, fixed or mobile.

The customer premises integrator (CPI) will have the task of devising gateway interfaces which will set up the call, connect to the appropriate service / access networks, and translate the signal into a format that will be recognized by such diverse home equipment as the computer workstation, multimedia creation center, home entertainment center, and security / energy management systems. Such a gateway will have to be dynamic and scalable, that is, capable of responding to a range of transmission / compression options and schemes, so that differently formatted services can be reconfigured on the fly. It must be able to manage the intelligent sorting and distribution of those services among diverse users within the home environment. Yet to be solved is how to account for usage, devise a service bill, and protect copyright. This will be no small matter. Some operators rely on metering of use and per-product purchases to determine charges; others rely on subscription contracts based on a package of services delivered or availability of on-demand services over time. An integrated accounting system is needed, one that seems cost-acceptable to users and proves profitable to content and service providers.

Summary

The home user stares out at a daunting array of competing products and services. These are positioned to reflect the strengths of their providers, many of which the average customer would love to sample from time to time. Unfortunately, what is offered comes in the form of all-or-nothing contracts on incompatible platforms. Most families would not be able to afford many of these even if they could arrive at a decision as to which one to choose. The challenge to the telecommunications industry is to integrate services and platforms and to work toward affordable, easy-to-use gateways that funnel information products and services into the customer premise and under the control of the home and small business user.

Useful Web Sites

@Home Network: *www.home.net*
The ADSL Forum: *www.adsl.com*
The ATM Forum: *www.atmforum.com*
Cellular Vision: *www.cellularvision.com*
Digital Audio-Visual Council: *www.davic.org*
Dudley Lab: *www.dudleylab.com*
Eon Corporation: *www.eon.com*
Federal Communications Commission: *www.fcc.gov*
The Institute for Telecommunications Studies, Ohio University: *www.tcomschool.ohiou.edu/itshome.htm*
International Telecommunication Union: *www.itu.ch*
ISR Global Telecom: *www.isrglobal.com*
MediaOne: *www.mediaone.com*
Network Management Forum: *www.nmf.org*
OptaPhone Systems: *www.optaphone.com*
Paradyne Corporation: *www.paradyne.com*
Road Runner: *www.excalibur-group.com*
Telecommunications Information Networking Architecture: *www.tinac.com*
Telegroup: *www.telegroup.com*
USA Global Link: *www.usagloballink.com*
The Veda Home Company: *www.veda-home.com/TelecomTom*

References

Bell, Trudy E., John A. Adam, and Sue J. Lowe. "Technology 1996: Communications." *IEEE Spectrum*. January 1996, pp. 30-41.

Boeke, Cynthia. "Surfing Among the Stars." *Via Satellite*. November 1997, pp. 28-42.

Cioffi, John M. "ADSL: The Beginning of the Information Superhighway." *1997 Annual Review of Communications*. Vol. 48, 1997, pp. 395-399.

Coleman, Price. "Making Sense of Set-Tops." *Broadcasting and Cable*. October 27, 1997, pp. 51-52.

___________. "Adelphia Plans Digital Blitz." *Broadcasting and Cable*. November 10, 1997, pp. 59-62

Conchon, Alain and Martine Lapierre. "Open Creation Environment for Operations and Services." *Alcatel Telecommunications Review*. 3rd Quarter 1997, pp. 232-237.

Cook, William J. "1997: A New Space Odyssey." *U.S. News and World Report*. March 3, 1997, pp. 45-52.

Cornu, J. "Networking: Paradox or Challenge." *Alcatel Telecommunications Review*. 1st Quarter 1997, pp. 2-5.

Edwards, Morris. "Gigabit Ethernet Takes Off." *Communications News*. July 1997, pp. 64-65.

_____________. "FCC Sets Auctions for LMDS, Upper 800 SMR Channels," *Wireless News*. August 7, 1997, pp.1-2.

Flannagan, Patrick. "The 10 Hottest Technologies in Telecom." *Telecommunications*. May 1997, pp. 25-32.

Gilder, George. "Sand, Glass, Air, and the Coming Bandwidth Tidal Wave." *1997 Annual Review of Communications*. Vol. 48, 1997, pp. 59-63.

Green, Heather. "All Roads Lead to the Web." *Business Week*. November 24, 1997, pp. 150-151.

Gross, Neil. "Clearing the Traffic Jams in Cyberspace." *Business Week*. November 17, 1997, p. 82.

Henry, Shannon. "End of Phone Companies as We Know Them?" *Washington Technology*. August 29, 1996, pp. 1, 94.

Holliday, Clifford B. "The Residential Gateway." *IEEE Spectrum*. May 1997, pp. 29-31.

Howes, Karen JP. "Satellites in the Age of Interactivity." *Via Satellite*. February 1996, pp. 114-126.

Kraemer, Joseph S. "Local Competition: The War of Many Against One (1996-2001)." *1996 Annual Review of Communications*. Vol. 49, 1996, pp. 325-341.

Lambert, Peter. "Inverse Multiplexing: ATM Cost Cure." *Inter@ctive Week*. November 3, 1997, pp. 33-36.

____________. "Microsoft's Cable Play: PC and TV, Too?" *Inter@ctive Week*. June 16, 1997, p. 8

Littleton, Cynthia. "Wireless Takes on L.A." *Broadcasting and Cable*. March 24, 1997, p. 28.

Machan, Dyan. "Broadcasting is Finished." *Forbes*. October 6, 1997, pp. 43-45.

Mack, Toni. "Communications: The Next Wave." *Forbes*. October 6, 1997, pp. 70-73.

McGarvey, Joe. "Cable Modems Draw Hopeful Crowd," *Inter@ctive Week*. June 9, 1997, p. 25.

Minoli, Dan. "Supporting Multimedia and Other Evolving Applications Using Broadband." *1997 Annual Review of Communications*. Vol. 48, 1997, pp. 729-741.

Morris, Andrew. "The Future Is Pegasus for TWC." *TV Technology*. June 19, 1997, pp. 14, 18.

Paradyne Corporation. *The DSL Source Book*. Largo, Florida, Paradyne Corporation, 1997.

Richards, Kristian. "Internet Via Satellite: The Teledesic Model." *Research Monograph*. Institute for Telecommunications Studies, Ohio University (www.tcomschool.ohiou.edu / its2.htm), 1997.

______________. "Internet Via Satellite: The Hughes DirecPC Model." *Research Monograph*. Institute for Telecommunications Studies, Ohio University (www.tcomschool.ohiou.edu / its2.htm), 1997.

Rooney, Lauren. "CAI First With Two-Way Digital Wireless." *TV Technology*. April 10, 1997, p. 24.

Saxen, Peter, and Ken Cavanaugh. "Rejuvenated ADSL Scores High." *1997 Annual Review of Communications*. Vol. 48, 1997, pp. 585-592.

Stewart, Alan. "The Battle for Bandwidth." *Communications News*. May 1997, pp. 36-41.

Tedesco, Richard. "Modems: Fast Bucks for Cable?" *Broadcasting and Cable*. August 11, 1997, pp. 43-48.

Waring, David L. "Customer-Premises Networks Come Home." *Bellcore Exchange*. Winter 1996-97, pp. 18-22.

____________. "Issues and Industry Activities in the United States on Broadband Customer Premises Networks." 22nd European Conference on Optical Communication. Oslo, 1996.

Wilson, Carol. "Access: The Next Generation (Technology Fuels Service Opportunity)." *Inter@ctive Week Supplement*. November 3, 1997, pp. S-2-12.

Zimmerman, George. "xDSL Solutions for High-Speed Connection Over Existing Twisted-Pair Copper." *1997 Annual Review of Communications*. Vol. 48, 1997, pp. 619-623.

Intranets and Corporate Use of the Internet

Caitlin M. Groom
Senior Consultant
Deloitte & Touche Consulting Group

Frank M. Groom, Ph.D.
Professor of Information and Comm. Sciences
Ball State University

Introduction

Corporations have discovered that the combination of using the public Internet and incorporating Internet tools on their private corporate data networks are efficient means to deliver information to a much wider range of interested parties. Corporations have interconnected their employees at headquarters and throughout the country with local and wide area networks for decades. Moreover, router-based TCP/IP networks have been popular for the past 15 years or more. Recently, the addition of Internet technology to a corporation's set of communications facilities has allowed companies to broaden the base of parties with whom they can communicate, while offering the advantage of presenting a common set of tools and protocols, which promise to be more efficiently supported.

The Components of an Intranet

An intranet can be defined as employing Internet technology over a corporate data network to access and provide information. The strength of a corporation employing an intranet is in the fact that common Internet technology, which is widely understood and deployed, is used to offer information over an internal data network to users. Further, these private intranets can be connected to the public Internet, providing a simple means for national and international information exchange. A firewall is normally placed at the intersection point between the internal intranet and the public Internet to protect against unauthorized user access to critical corporate assets.

The network of the intranet is implemented with TCP/IP software operating over routed or switched networks. The universal vehicle by which users access these networks when reconstituted as intranets is a browser, such as Netscape's Navigator or Communicator or Microsoft's Internet Explorer. The data which users seek to access over an intranet is stored as documents on a server in hypertext markup language format. Images, GIF and JPEG pictures, forms, templates, and links to other pages and servers are also employed to form the Web server information database. Since information abounds in corporate legacy systems, as well as more distributed departmental, work group, and individual file systems, data pump software has been developed to extract data which is stored in its original format, frequently as fixed fielded rows amalgamated into tables for the relational database structure. The extracted data is then reformatted into an HTML format and reconstituted as a document. The document is then placed on one of the corporate Web servers.

Many systems, such as Lotus Notes or client/server structured systems require formal "client" software on a user's PC to interface with a server-based application software, as well as remote procedure call routines to interface to the network protocols. JAVA applets have been developed to be placed on user PCs so that they can augment the capability of a user's browser. As a result, the user can access and "browse" both Web servers and more traditional systems regardless of how the desired data is stored and formatted.

When the corporation wishes to make nonsensitive data available to the general public, customers, suppliers, distributors, shippers, or remote employees, it has the option to place the data on the Internet side of the firewall. In this situation an unprotected or password-controlled Web server is connected directly to the Internet as a selectable designation server with a standard IP address. Such an external Web server is considered an extranet, since the server and its data are "extra," or outside the firewall. The protected information remains inside the firewall on the intranet with its "intra" network, servers, data, and users.

Structure of the Internet

The public Internet is constructed as a distributed network composed of four basic components. The core network is composed of a set of high-speed backbone links interconnecting the four access points to the network. These access points are termed network access points (NAPs) and are placed in San Francisco, Chicago, New York, and Washington, DC. Access to these NAPs is achieved either directly, for a limited set of users such as the national supercomputing centers, or through DS-3 45-Mbps links from a carrier's (either Sprint's or MCI's) local point of presence (POP). A company may connect directly to such a carrier POP for Internet access. However, individuals or small organizations connect by means of subscribing to Internet access service from an Internet service provider (ISP), which connects to the carrier

FIGURE 1

The Structure of the Internet

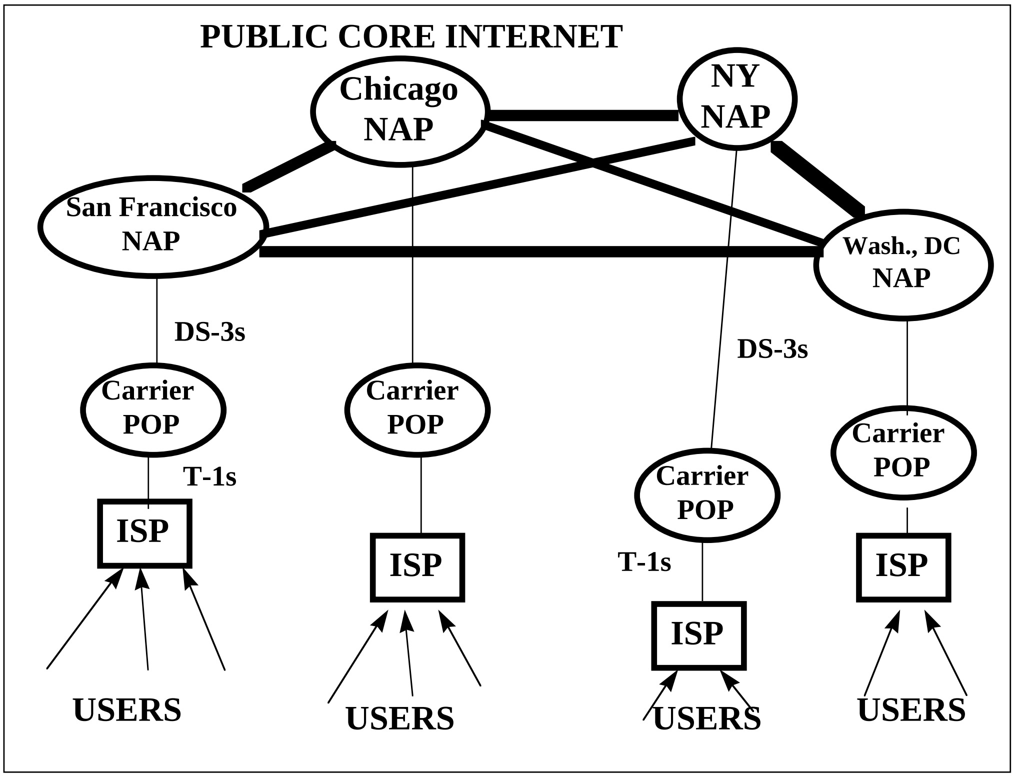

and on to the NAPs for them. Internet directories are maintained in local domain name servers (DNSs) at the ISP or local corporate network and are pointed to by ten national DNSs. *Figure 1* portrays the Internet structure in the United States.

Although the core interconnection linkage of the Internet is physically in the United States, Europe and Southeast Asia have developed country-based sub-Internets, which are linked to the U.S. Internet. Europe has developed its own set of inter-links to interconnect the Internet networks of individual countries in Europe and to connect to adjacent countries. Southeast Asia tends to have country Internets that are directly connected by a carrier link to the U.S.–based Internet. This is portrayed in *Figure 2*.

The Advantage of Deploying Information with an Intranet

The advantage of deploying information on an intranet versus the traditional way in which data has been made available in standard database systems on a corporate data network is the employment of common user software, a common method to access data, and finally a common method to store data and make it available to users, both at internal locations and throughout the world. This technology is well known to users and technicians. It can be managed with standard tools and deployed with common user and server software components. Data that the corporation wishes to be made generally accessible can be placed on a Web server, as an unprotected extranet Web server or behind a firewall on a more protected intranet. Data that must be private and protected remains in secure legacy system databases.

Conditions Affecting the Internet/Intranet/Extranet Decision

In deciding how to deploy an intranet or extranet, the user community and their location must be determined, as well as identifying the networks to which they have access, the software and hardware tools they employ, and the level of protection required for the data being offered.

Information Users

The first decision in building a corporate intranet is to determine the complete range of users who will make use of the network to access information. It is essential to know where they are located and what type of network connection is available. If they are in a work environment, do they have a local network connection to an edge router through which they can access the Internet? Alternatively, it may be possible to determine whither they have access to the corporate wide area network, either directly or in a dial-up mode.

It is possible to target general employees, professionals, corporate managers, and executives. They may be given access whether they are in their corporate office, in a branch office, working at home, at a customer location, in a hotel, or using mobile communications such as with cellular or personal communications service (PCS) connection.

FIGURE 2

Asian, American, and European Internet Linkage

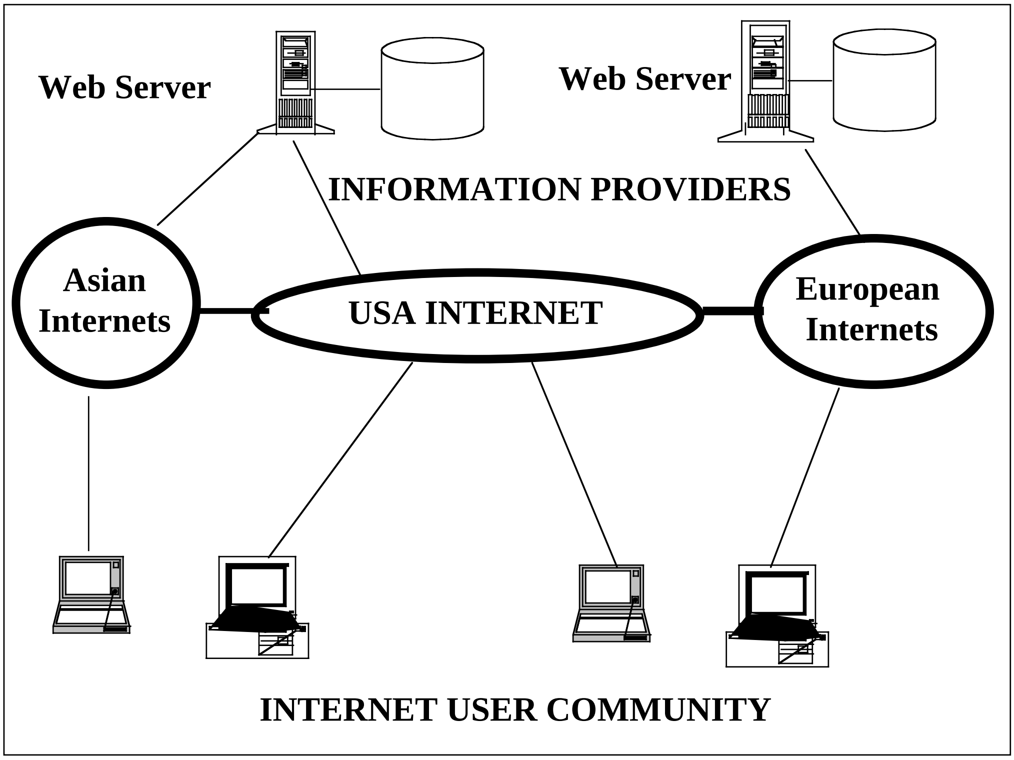

However, Internet communication is not limited to corporate employees. The complete supply chain, including suppliers, assemblers, manufacturers, warehouse and shipping companies as well as customers, is becoming linked over the Internet, by dial-up, or through private network connection to a corporate intranet. Moreover, by placing information on Web servers outside the internal corporate intranet, on extranet servers directly connected to the public Internet, a company can attract, inform, and educate a wide range of potential new customers both nationally and worldwide.

Depending on the location of these individuals and entities and their local communications capability, decisions can be made as to the most appropriate vehicle for their access to the desired corporate information, whether to place such an information server inside or outside a protective firewall. *Figure 3* portrays such a central information location, accessed by a mixed set of users who are located throughout the United States and Canada.

Common Software

The advantage of employing a common browser, such as Netscape's Navigator, is that the same access software can be deployed on virtually all user processing platforms. It can operate on all Intel 386, 486, Pentium, Pentium Pro, or Pentium II PCs, regardless of whether they employ Win95, NT, or Win 3.1 operating systems. Versions are also available for Sun, MAC, and PowerPC computers. These user PCs, equipped with a browser, can access a remote server by dial-up connection with a modem over a telephone connection, over a local area network (LAN), over a wide area network (WAN), or through any other network connection. Most users are familiar with these browsers already, and many support technicians have already experimented with Web server software. With the increasing familiarity many have, through creating their own Web pages, with the Web server CGI and PERL Languages, HTML formatted information, and the creation of "hot links" between Web components, a large and growing supply of potential support personnel is available.

Firewall Protection

Corporations must protect their corporate information from unauthorized intruders. Although they might choose to interconnect their internal intranet to the public Internet they must still protect certain offered data as well as the bulk of information used to operate the company. *Figure 4* portrays two means for user to access corporate data. Some users are gaining access by means of connecting to the Internet through modems and telephone connection. The Internet then connects them to the corporate intranet where the desired information resides. Corporate users, on the other hand, might gain access by means of corporate WANs connecting to the central corporate data network that supports the corporate intranet. A firewall allows authorized traffic flow between these two networks—the corporate intranet and the public Internet.

Since traffic from the public Internet can, accidentally or deliberately, be malicious, routing and server products have

FIGURE 3

Users Throughout North America Accessing a Central Information Site

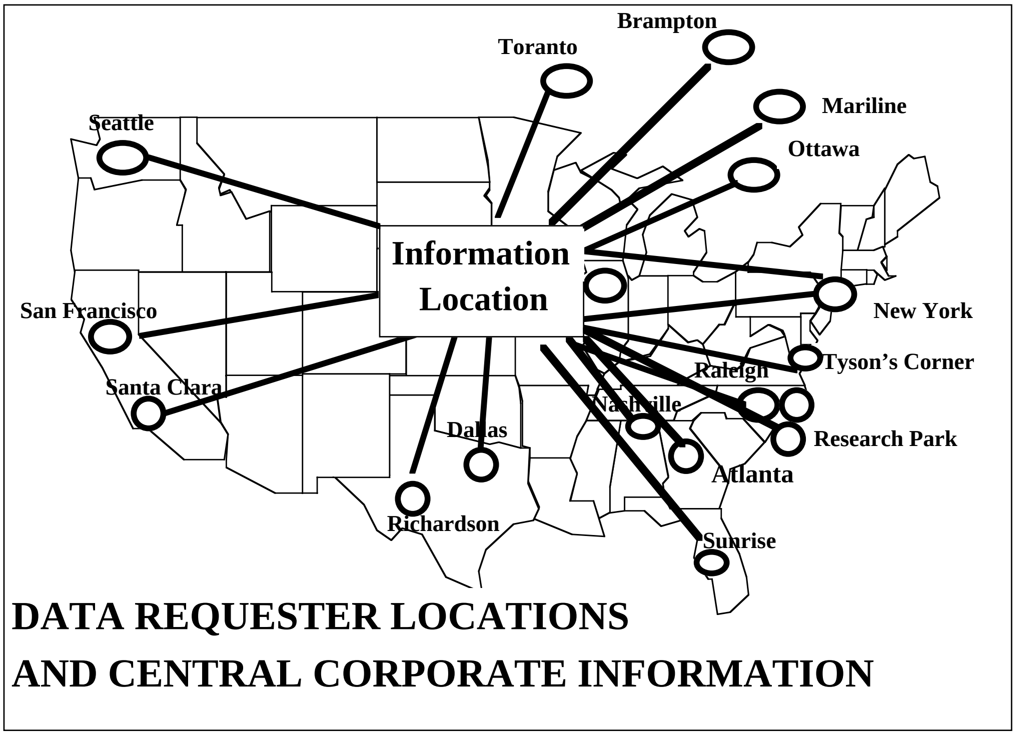

FIGURE 4

Alternate Means to Connect to the Internal Corporate Intranet

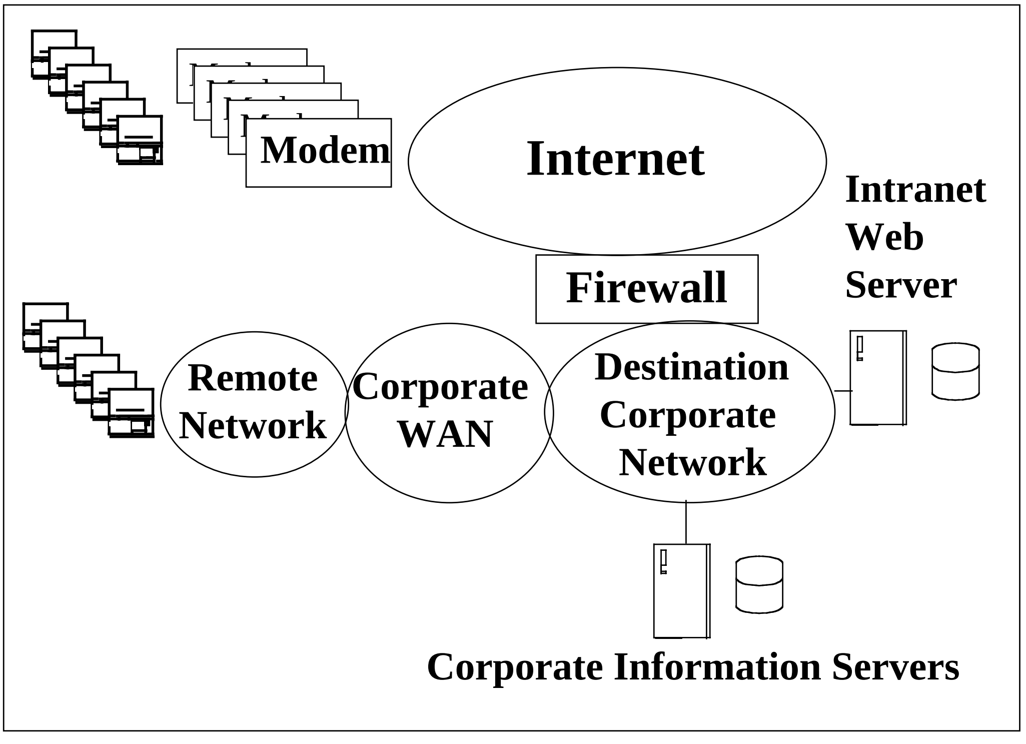

FIGURE 5

A Selection of Firewall Products

FIREWALLS		
	Integralis	NetGate
Blackhole	Interceptor	Netpartner
BorderWare	Inter-Ceeptor	Netra Server
Brimstone	ANS Interlock	NetSP IBM
CENTRI Secure	Internet Secure Rtr	Network-1
Cyberguard	Turnstyle TIM	Novix
DBF	IRX Router	PIX
Eagle	IwareNetGate	PORTUS
ExFilter	KarlBridge	Quiotix
Firewall-1	Mazama	SEAL DEC
Guantlet - TIS	MIDnet SecurIT	SecurityGate DEC
GFX-94	NetCS	Sidewinder

been developed to provide a screening and authorization process, commonly known as a firewall. *Figure 5* presents a list of some of the current firewall products.

Connectivity and Existing Corporate Data

Initially, intranets were considered to only include browser access through TCP/IP networks to Web server-stored data. Data pump software was required to move legacy data from a legacy system database operating on a mainframe computer in a corporate data center to the desired Web server. These legacy systems were historically written in the COBOL language, and stored their data in either IBM's hierarchical IMS or relational DB2 databases. Sometimes Oracle's database system or Sybase's system for smaller computers might be employed. The network that connected the user to these systems was predominately IBM's point-to-point SNA network. In the late 1980s and early 1990s many such systems were converted to a client/server structure with a user client software package, a remote procedure call software module to interface with network software, a TCP/IP networking package, and Unix servers with C or C++ servers running the applications and maintaining the stored data. Such systems were usually distributed to locations close to the user community.

In the mid-1990s a German company began offering the SAP system for building integrated corporate systems, offering a set of pre-built processing modules and a programming language for the creation of customized processing routines to extend the pre-built modules. SAP systems can operate with many diverse storage products, but most frequently are implemented with either Oracle or Sybase's relational database management system. Furthermore, as corporations downsized, individual workgroups implemented their own information sharing systems, with Lotus Notes being the most popular product employed. The legacy systems, the client/server systems, the SAP systems, and the Lotus Notes systems each required their own specific "client" software package to run on a user's PC in order to access data on the appropriate server system and its database. On the other hand, with the intranet approach the Web server expects the user to employ a browser package to access Web-based information. This diverse mix of information systems and databases that a remote user might wish to access are portrayed in *Figure 6*.

Since all required data may not be downloaded to the corporate Web server for Internet or intranet access, JAVA applets have been created to present fill-in-the-blank forms for upload to the non-Web systems, simulating their anticipated client format for requesting information.

Building the Corporate Intranet-Web Server Software

Traditionally, UNIX-based servers were used to offer information over the Internet. However, Microsoft's Windows NT has become increasingly popular for a wide range of corporate information delivery systems. Products to construct and operate a Web server are specific to the operating system deployed. Since UNIX has been available for so many decades, many Web server software products are available for these UNIX-

FIGURE 6

An Array of Information Systems a User Might Need to Access

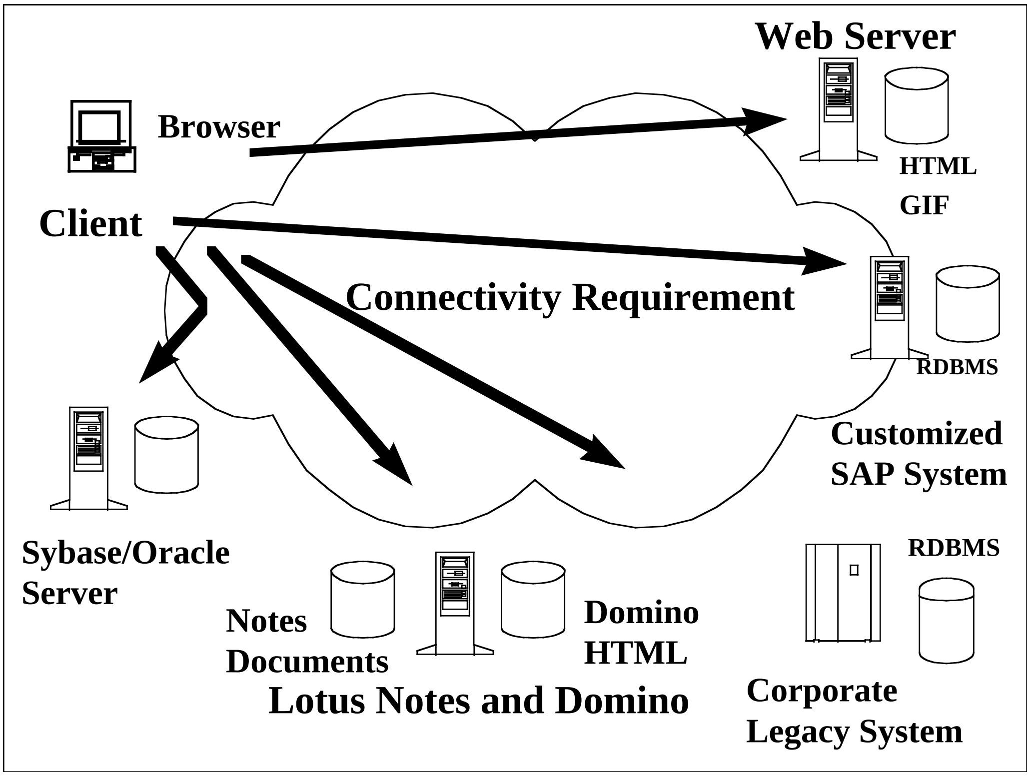

based servers. Apache is one of the most commonly deployed with the added attraction that it is free. Other products, such as Netscape's Enterprise server are modestly priced at $995 per server. A sample set of Web server construction software for UNIX systems is presented in *Figure 7*.

Microsoft's Windows NT is being deployed more and more often for corporate application and information server processes. Many Web servers are now being constructed for machines using the NT operating system. *Figure 8* presents an example of Web Server products for use in building a Web server under the Windows NT operating system.

Moving Data from Legacy Databases to Web Server Databases

Over the past 35 years corporations have placed and managed the core information of the company in large data processing systems, frequently referred to today as "legacy" systems. These systems may store data in hierarchical databases, relational databases, flat files, or more recently in document files such as Lotus Notes. Data is stored in fielded records for databases and files, and in a free-form document format for Lotus Notes. On the other hand, information is stored in a radically different HTML format for Web servers. Thus, a reformatting process is required when data is moved from a legacy system to a Web-based system. Furthermore, the data must be located and then accessed on the legacy storage systems, extracted, reformatted, and then entered into the Web server storage—a series of intricate processes. Software packages have been developed to assist in this extraction and reformatting process. IBM has developed its MQ Series for the extraction process. *Figure 9* presents a set of such products, the offering company, and the price where available, that exist to perform this extraction and conversion process. It further presents the types of formats and databases each can access.

Finding Required Data Stored on Web Servers

Once a user decides what data is required, he or she faces the problem of finding where it is located, since corporations might deploy data on many Web servers at many locations. Further, it is often efficient to store different kinds of data on separate servers, with the data interconnected by HTML "hot link" code. In order to start this process, a user needs find and access the appropriate location of a core set of data that is desired. To accomplish this search for the appropriate server location, the user first accesses the internal corporate search engine software and provides the appropriate data names, which can be used as key words to find the stored data. Similar to finding information on the Internet, these search engines inventory all the data elements, names, and locations. Accessing the search engine provides a means to find and then link to a starting set of core data that is requested. From that starting point links to other data can be followed until the user's information requirement is satisfied. Many companies offer these search engines, with the most successful being DEC's Alta Vista, a smaller version of DEC's popular offering

FIGURE 7

A Selection of Web Server Software for Unix Servers

WEB SERVER SOFTWARE- UNIX		
Apache	Apache	Free
Common Lisp-HTTP	MIT	Free
Cosmos Web Server	RIS Tech.	$2k-$20,000
ExpressO HTTP Srvr	Captal City	$69-$195
Internet Connection Secure service	IBM	$295
InToto	Computat'lLogic	$1950
Jigsaw	W3 Consortium	Free
mhttpd	Universal Access	Free
NCSA HTTPD	NCSA	Free
Enterprise Server	Netscape	$995

FIGURE 8

A Selection of Web Server Software for Windows NT

WEB SERVER SOFTWARE- WINDOWS NT		
Basic Web	Info Dimensions	$2500
Oracle Webserver	Oracle	$2495
Wildcat Interactive Net	Mustang	$249-$3999
ZBServer	ZBSoft	$25-$69
ISYS Web Server	ISYS/Odyssey	?
mhttpd	Universal Access	Free
Spinnaker Web Server	Searchlight	$495

FIGURE 9

A Selection of Data Pump Products

DATA PUMPS			
NotesPump	IBM Lotus Notes	RDBMS	$7,995
Data Integrator	Sentinel Software	Notes, Informix, ASCII	$10,000
Replic-Action	Casahl Technology	RDBMS, Notes Legacy DBs	
InfoPump	Platinum Technology	RDBMS, Notes PC DBs, Flat Files	
SQL Pump	Softquest	RDBMS	$995/5 $3995NL
NotrixComposer	Percussion Software	RDBMS	
Datalink R3	Brainstorm Tech.	Rdbms, Notes, ASCII	

for the public Internet. *Figure 10* presents a sample of available corporate search engines to be used in locating and accessing Web server-stored information components.

Push Technology

In contrast to a user taking the action to access data, many companies are taking the reverse approach of scheduling information to be automatically downloaded to a set of users at appropriate, prearranged times. PointCast developed the first system to automatically download requested information to registered clients over the Internet and remains the most popular provider of that technology. Many companies now use push technology to automatically send reports, announcements, and new products specifications to employees and customers over the Internet. A number of other competitors have emerged offering alternative products. An example set of such vendors is:

- PointCast with its I-Server products priced at $995 for an unlimited number of users

- Intermind with its Communicator products priced at $995 for 100 users, and an additional $195 for each copy of its user client software

- Marimba with its Constanet product costing $1,000 for 100 users, with its client software being free

- Wayfarer Communications with its Incisa and Databridge access software costing $5,000 for 100 users while its client software is free

- Backweb Technologies used with Lotus Notes costing $10,500 for 300 users and $10–$35 for each additional user

- DataChannel with its channel manger with open connect systems products costing $6,900 for 100 users and between $30 an $60 for each additional user

- Netscape offers its Constellation system while Microsoft offers its Active Desktop

Examples of Internet/Intranet Usage

Many companies are now using the Internet as a routine means to convey information. The following presents a sample of the variety of approaches corporations are employing to make use of intranets in combination with the public Internet to deal with the wide range of individuals and business entities that might need access to a corporation's information.

Extranet Information Servers

Companies such as Intel, Microsoft, Motorola, Lucent, Cisco, and Sun have discovered that they can reach a vast audience if they place information in a public Web server and allow people to access it over the Internet. They post employment opportunities and anyone throughout the world can submit

FIGURE 10

Corporate Search Engines

INTERNAL CORPORATE SEARCH ENGINES

Company	Product	SQL	Push	Html Links	No-Web Doc Types	Price
DEC	AltaVista			Y	Y	$15,995-$99,995
Folio	Site Director			N	Y	$9,995
InfoSeek	UltraSeek			Y	Y	$995
Lycos	Int. Spyder			Y	Y	$7,500
Odyssey	ISYS:Web			N	Y	$5,200
Maxum	Phantom			Y	N	$395
Microsoft	Index Svr	Y		Y	Y	
Netscape	Compass Svr		Y	Y	Y	$1,200- 50 users
Quadralay	Web Works Search			N	N	$495
Verity	IntelliServ		Y	Y	Y	$2,995- 5 users $8,990-25 users

a resume and application for a job. Furthermore, when Lucent technologies separated from AT&T, they were able to post information to employees, suppliers, and customers about the separation sequence. They could also let employees know where they needed an increased number of workers and employees could act on their own behalf in applying for internal transfers.

Cisco has used the Internet to deliver the complete set of their technical manuals and specifications to technical support people throughout the world. When configuring a router to handle frame relay, both the customer technician and the local Cisco support person can access the specific instructions, which are stored and made available at the central Cisco Web site, jointly. Furthermore, a vast number of "white papers" and design suggestions are available for all forms of router network deployment and design.

Sun has done much the same, with the unique example of placing the Java language and examples of Java code for anyone to access and employ. Detailed technical "white papers" are available for download. *Figure 11* presents this concept of an extranet Web site for delivering such an array of information and receiving responses where appropriate.

Specific Corporate Intranet/Internet Applications

A number of companies have set up intranet sites behind firewalls for delivery to external Internet users. Each has a unique problem to address with this technology and has a slightly different approach in delivering required information. Among these companies are:

Fruit of the Loom

Fruit of the Loom has employees and distributors at many locations. They use the Internet to distribute production reports, inventory analysis, and projections, as well as required employee data. Rather than having each employee access the central Windows NT server that maintains this information, Fruit of the Loom uses PointCast's push software technology to send the information to the user's browser. The information server is placed behind a firewall on a Fruit of the Loom intranet. They download data from legacy systems using PointCast I-Server auto-retrieve software, format it for distribution, and store it in an information database on the Windows NT server operating as a staging device, ready to then send the information when the appropriate time arrives. *Figure 12* portrays Fruit of the Loom's push technology information delivery scheme.

McDonnell Douglas

McDonnell Douglas maintains communications with over 700 suppliers and 7,500 employees throughout the world by means of the Internet. They have constructed a corporate intranet behind a firewall. Inside the firewall they have constructed a Lotus Notes server with which they communicate through the firewall, over the Internet, to this wide array of users. The material to be distributed comes from two sources. Corporate information from legacy data systems is downloaded to the Lotus server by IBM's MQ Series data pump. Additional information is supplied over the internal corporate intranet by users employing either a browser or Lotus Notes client software. Technical documentation and

FIGURE 11

Delivering Information by Means of a Corporate Extranet

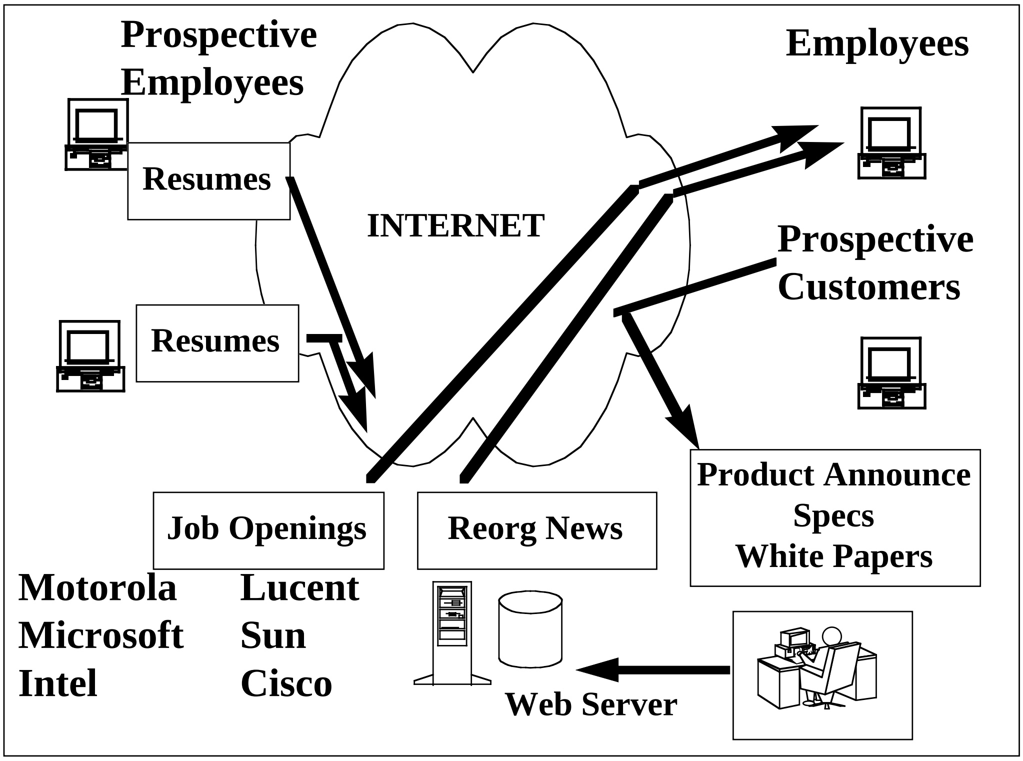

FIGURE 12

Fruit of the Loom's Push Technology Information Delivery System

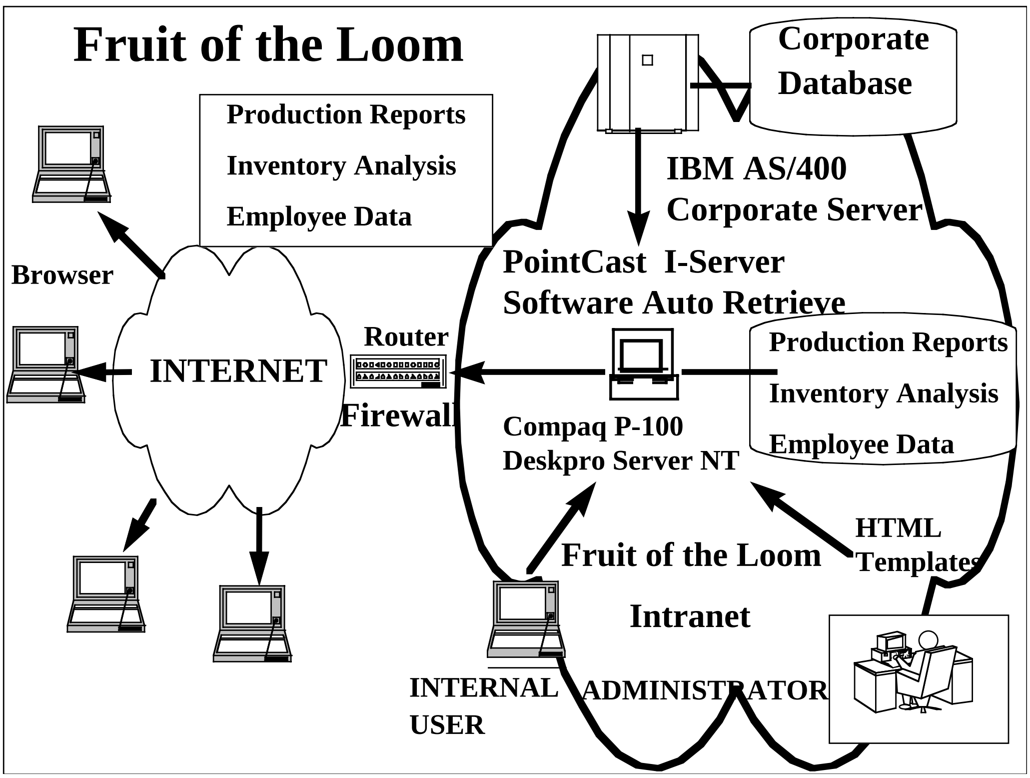

FIGURE 13

McDonnell Douglas's Intranet/Internet Information Delivery System

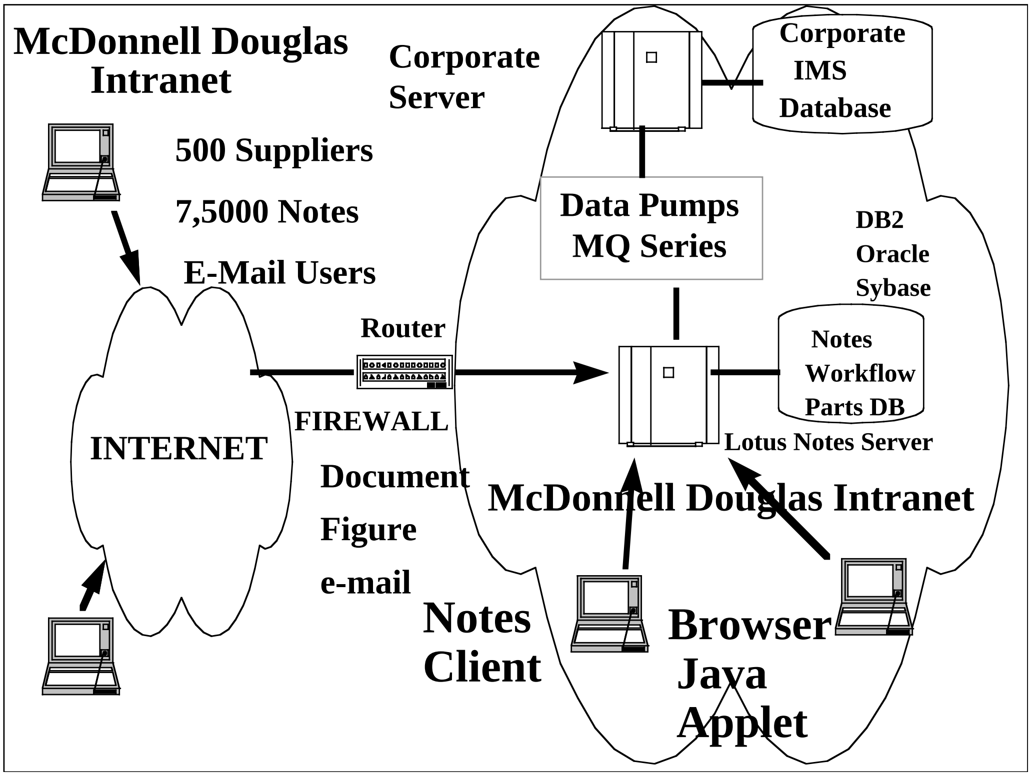

specifications, reports, figures, and e-mail are then sent from the Notes server through the Internet to browser-equipped users. *Figure 13* presents McDonnell Douglas's Notes-based intranet / Internet information delivery scheme.

Fidelity Customer Trading via Internet Order Submission

Fidelity has constructed a double firewall-protected Internet security trading system. Fidelity allows customers, using Netscape's browser with secure socket layer (SSL) implementation, to access a Fidelity Web server over the Internet. The transaction is encrypted, and only authorized customers are allowed through the initial firewall. The customer then, using a downloaded form, submits a buy or sell order. Fidelity's Web server then sends the customer's order request to a transaction server which, using client / server processing, submits an order for the customer through another internal firewall to a banking system server. The Fidelity banking system server then submits the trade request through its normal client / server process to the security trading computer and updates the account management system of the trade details. *Figure 14* portrays Fidelity's Internet security trading system.

Streamline Grocery

Streamline has developed a useful grocery service for families with two working parents, one-parent families, and traveling workers. Streamline places a refrigerator in the garage of the customer with which it has arranged authorized access. The customer can then place a grocery order through the Internet to Streamline. The company then transmits the order to the closest warehouse / delivery location which, on the requested date, delivers the ordered goods to the customer's refrigerator. When the customer arrives home, the refrigerated food awaits preparation for the night's meal. A follow-up invoice and bill are then submitted to the customer over the Internet. *Figure 15* portrays the Streamline system.

Internet Sales

Using the Internet as a universal sales vehicle has captivated the imagination of the press, but has been disappointing for all but a few. Among the more successful on-line vendors are the following:

- *Amazon.com*—Using Netscape's Secure Socket (SSL) encrypted HTML transactions, customers can place orders for books from a universal list stored on a DEC Alpha Server under an Oracle 7.2 database. Over 1.1 million titles are available for selection. Amazon.com ships from a warehouse maintained by Ingram Book Company, the nation's largest distributor of books.

- *Barnes and Noble, Borders, and Simon & Schuster*—Due to Amazon.com's success, the other major national book companies have created their own version for Internet book ordering, employing essentially the same technology as Amazon.

- *CD Now and CD Universe* —These companies have become the standard for CD ordering over the Internet,

FIGURE 14

Fidelity's Internet Customer Trading System

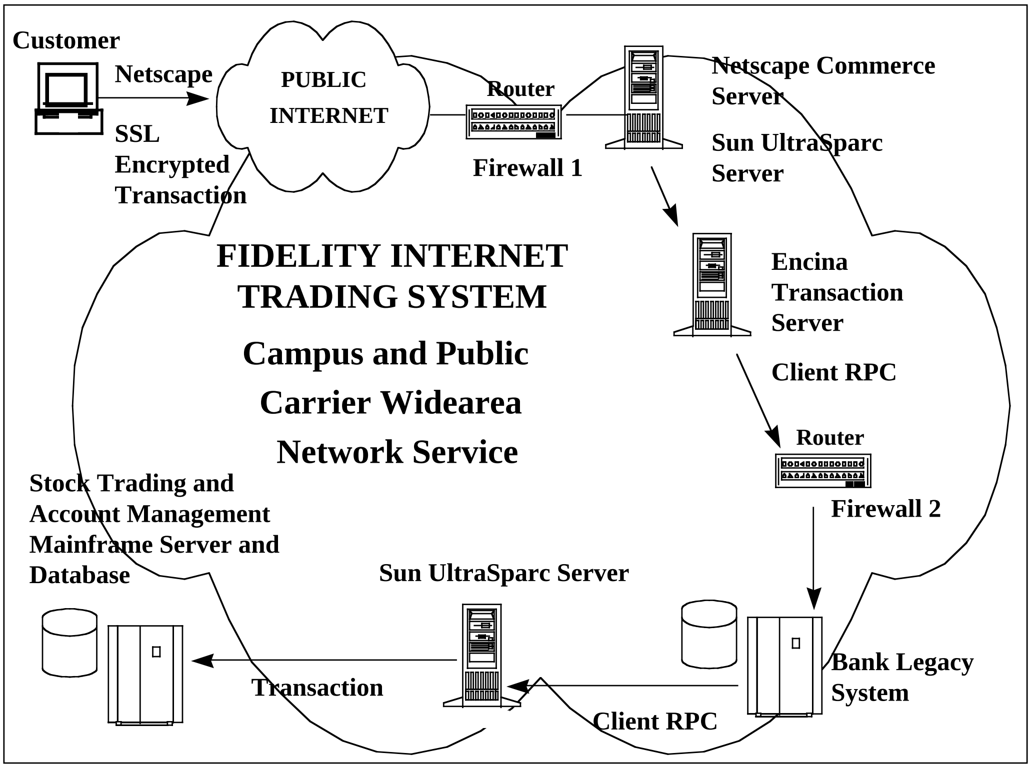

FIGURE 15

The Streamline Internet Grocery Delivery System

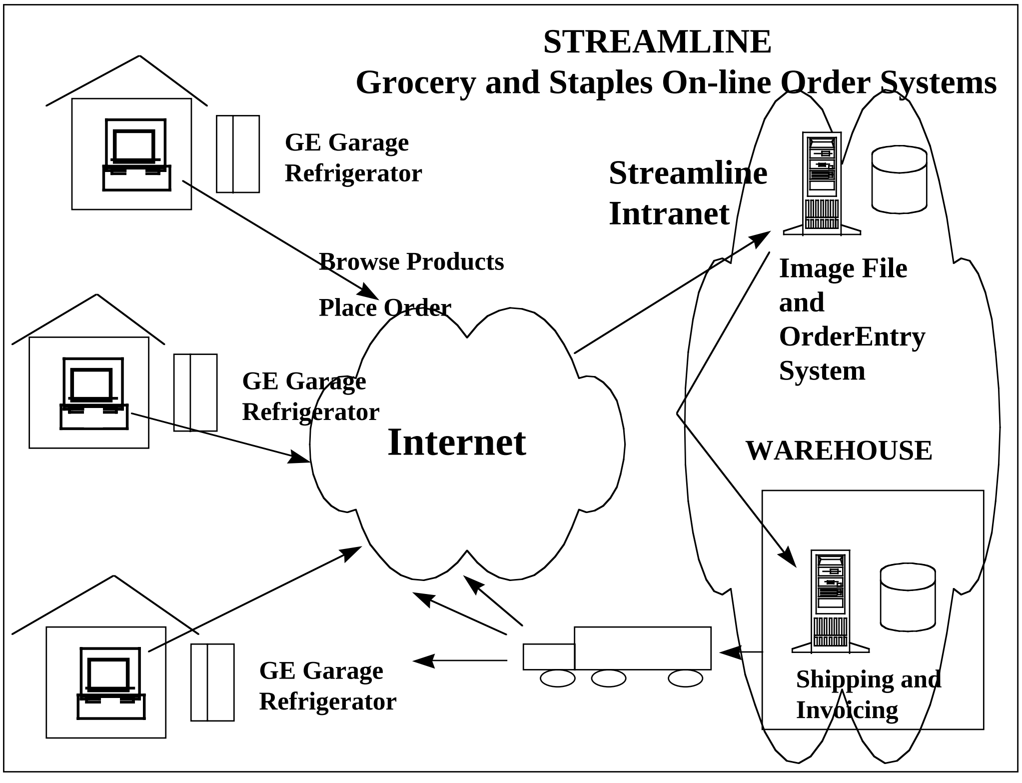

with CD Universe having the reputation for finding hard to get or out of print CDs.

- *Virtual Vineyards*—With music and books available on-line, it was only a small step to add the capability of ordering a good bottle of wine, cheese, and gourmet food for delivery to supplement a quiet evening of reading and music. Virtual Vineyards has become the prime Internet vehicle for such service. A good cigar must be not far behind.

Conclusion

Intranets provide a safe vehicle to construct an internal network using standard Internet technology. The construction of Web servers follows a widely understood format and technology. User access by means of browsers is simple and easily understood by most potential users. Search engines, such as Alta Vista, are commonly used on the Internet and transferable to searching for data on corporate intranet Web Servers. Since firewalls isolate the corporate intranet from the external public Internet, a relatively safe environment can be created while offering the simplicity of Internet technology to corporate users, developers, and administrators. This approach is proving a simple and inexpensive means for corporate communications and information delivery to a widespread user community while also supporting simultaneous data access and communication with an internal corporate employee body. As more sophisticated storage and access technology are developed based upon Web technology, corporate usage will continue to find the Internet and Internet technology a strategic vehicle for corporate communication.

Making Money with the Internet

Mike Heller
Senior Manager, Service Provider Line of Business
Cisco Systems

The Internet has exploded in all directions over the last few years. It is difficult to pass one day without hearing or reading something about the past, present, or future of the Internet. At this point, there are more than 100,000 connected networks with 10 million host computers and 76,000 Web servers serving 35 million users in 186 countries. The Internet may be one of the most amazing success stories in history. Not only is it big today, but there is an 85 percent annual growth in networks, a 300 percent annual growth in traffic, and a 2,400 percent growth in Web servers. There are already over 42 million home pages on the World Wide Web (WWW), and the estimated 130 million users in 1998 should jump to 200 million by the year 2000. Finally, over 65,000 businesses are connected, showing a growth rate of 14 percent per month. This means that bandwidth requirements are doubling every eleven months. Is this just a fad? Will the Internet explosion die away like so many other popular cultural phenomena? All evidence shows that the answer to these questions is no. This is a major turning point in telecommunications and the world economy. This paper will examine the Internet marketplace and the opportunities therein and outline some strategies for service provider success in this phenomenal area.

The Explosion

Are Internet opportunities equal for all service providers? A study by the U.S. Department of Commerce and Forrester Research, Inc., provides a glimpse at the answer. *Figure 1* shows data from companies ranging from over one thousand employees to those with 100 or less. These companies were asked, as a part of a market-research study, whether they were currently connected to the Internet and, if not, whether they intended to be, and when. The percentage of the largest companies intending to be Internet-capable by the year 2000 approaches 100 percent, but only about 60 percent of the mid-sized companies and only 25 percent of the smallest companies intend to be connected to the Internet by the year 2000. Given the incredible opportunities seen by larger companies, why do the smaller companies not share the same view? What do these figures suggest about the equality of opportunity?

There have been a variety of surveys examining small and medium-sized businesses. Results indicate that the major problem causing this disparity is that smaller companies do not see the value of the Internet. They have not studied Internet technology, nor have they carefully evaluated the business proposition and found no value to their organizations. Rather, they do not understand the benefits. This fact creates a marvelous business opportunity for service providers who are in a position to provide expertise and a complete solution perspective.

Another interesting area of growth has been with intranets. An intranet is an internal network within an organization that uses Internet technologies. The Internet is, in part, a public

FIGURE 1
The Internet Explosion

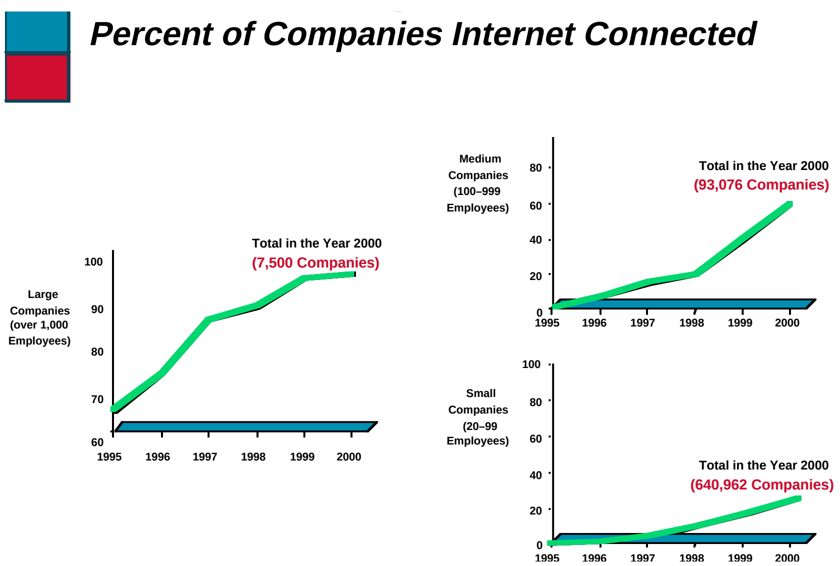

FIGURE 2

The Intranet Explosion

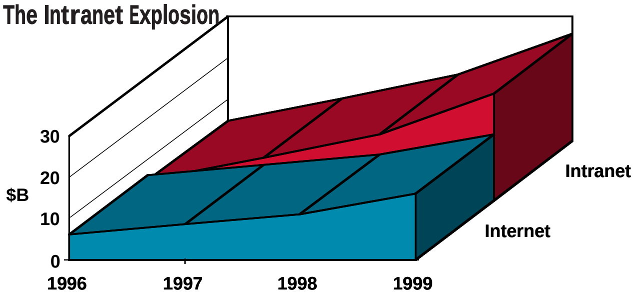

Zona Research Inc., Internet / Intranet Technologies and Services Market

network through which people communicate with those in other companies, organizations, and universities, while an intranet is a closed user group including the employees of a company and some selected customers or partners. *Figure 2* shows results gathered by Zona Research, Inc. that suggest that intranets will produce a higher amount of business than the Internet itself.

Not only is the volume of users growing at an amazing rate, but also the types of Internet activities have changed strikingly. When the Internet first began 25 years ago, the predominant applications were e-mail and file transfers of text or binary data. At present, there are not many text-only Web pages—most have graphics. The use of graphics means that more bits are being transferred across the network per a given user's activity. There are now over a thousand radio and television stations "broadcasting" over the Internet. These activities consume far more network resources than a simple e-mail message. This means that, for network service providers, the demands in capacity are increasing at the compounded volumes shown in *Figure 3*. Networks are growing at an incredible rate.

The Internet Opportunity

The Internet represents the single greatest business opportunity that telecommunications companies and service providers will ever have. A service provider that does not involve itself in this trend is making a potentially fatal mistake.

There are several reasons that the Internet is so successful. First, the graphical user interface (GUI) found in Web browsers is akin to a Macintosh interface or Windows interface. It has created a convergence of technologies. Second, personal computer and electronic costs have decreased significantly in the last few years, making the Internet available to far more people than was originally imagined. Due in part to this trend, the Internet has become a de facto standard.

Another reason for the huge success of the Internet is the creation of the World Wide Web. This is a relatively recent addition as compared to the Internet with its earlier applications of e-mail and file transfers, but it has become the most active area by far. Businesses and organizations can find the most immediate value from the WWW. E-mail applications are interesting and useful, but it is now possible to advertise, sell, and realize a profit through the Web.

A final cause of the Internet's success can be thought of as vested interest. Not more than two or three years ago, Bill Gates of the Microsoft Corporation stated publicly that he did not understand why the Internet was drawing so much focus. This is a view he no longer holds. Microsoft and almost every other large company now have a vested interest in the Internet, solidifying its position for the foreseeable future. Not only will the Internet continue to be supported, but also it is in the best interest of companies such as Microsoft to see that problems are fixed and that it remains as accessible and user-friendly as possible.

The Internet's time has come. It draws heavily on the computer evolution that has been in process since the 1950s, spanning the mainframe era, the minicomputer era, the PC era, and now the client-server or distributed-computing era (see *Figure 4*). This current era could be seen as the network-computing era. In mainframes, minicomputers, and personal computers, everything was happening in one computer, whether big or small. Today, it can take at least two comput-

FIGURE 3

The Net on the Net

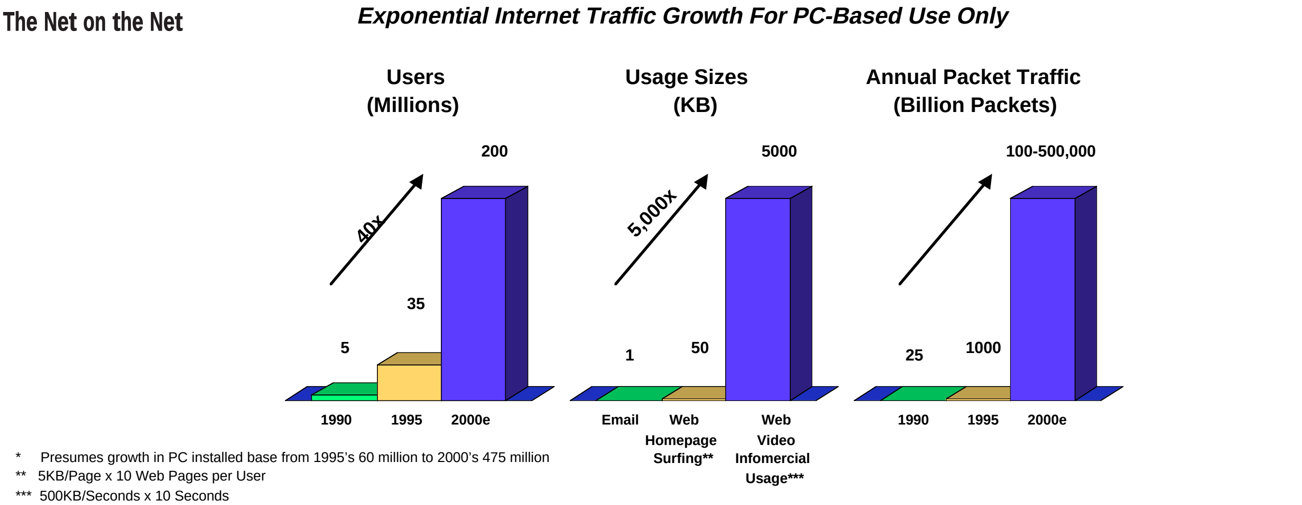

Source: IDC, Zona Research, Literature Searches, Team Analysis.

FIGURE 4

An Idea Whose Time Has Come

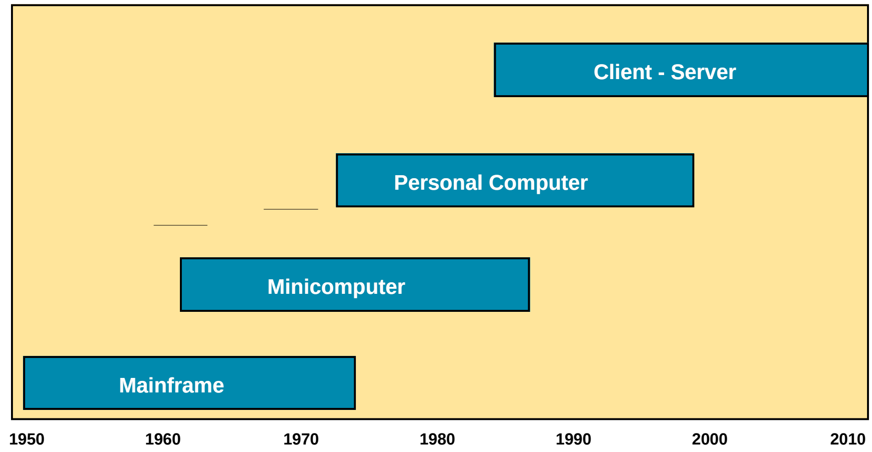

ers for something to happen. The Internet requires a person with a PC, Macintosh, or workstation, and another computer elsewhere. The presence of a network between the two is necessary for the application to work. This is good news for network providers.

The Service Provider Opportunity

Figure 1 indicates that many medium- and small-sized companies do not or are not planning to access the Internet, primarily because they do not understand its value for them. Large companies make up a relatively small percentage of total potential users. There are more medium-sized businesses than large ones, and more small businesses than medium-sized ones. Residential users make up the largest group. A Fortune 500 company with its own information-technology staff will not need an outside source to help it take advantage of the Internet opportunity. The medium and smaller users, however, provide an excellent possibility for service providers (see *Figure 5*). These users do not know what to do, how to take advantage of the opportunity, or even exactly what the opportunity is. This is the kind of information that service providers are in the ideal position to bring. They can provide service on demand, explain how these services can work to the benefit of users of any size, and, ultimately, gain these users' business and even loyalty.

Strategies for Success

For a service provider, the place to start is providing Internet access. There are, however, many other things that can be done, including the following:

- Internet access service
- Internet applications
- enhanced Internet applications
- customer premises equipment
- server hosting services
- security services
- content development services
- corporate remote access services

These companies should think in terms of being solution providers, not just access providers. Service providers can

FIGURE 5

The Service Provider Opportunity

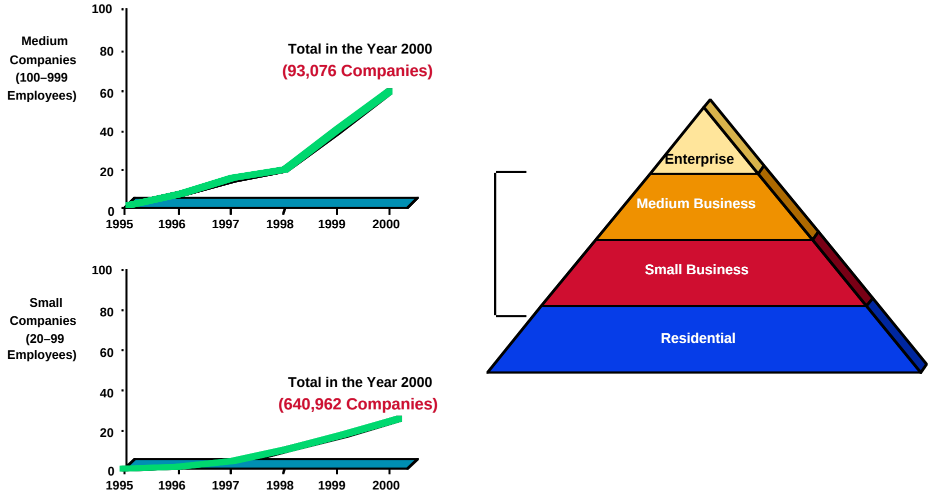

FIGURE 6

Access Environment

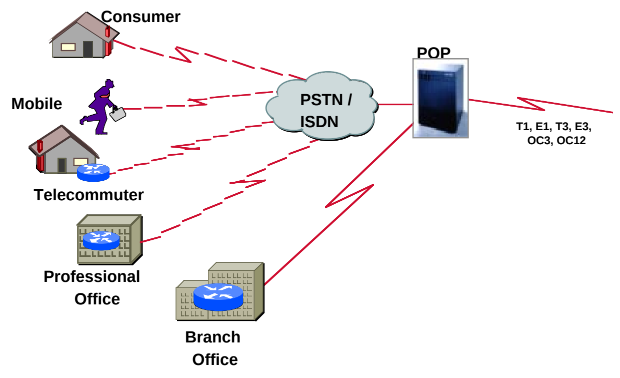

also branch into Internet applications such as e-mail. Enhanced Internet applications are another promising area. Combining computer systems with networking can provide very interesting and very powerful business advantages.

Providers could even sell customer-premises equipment. Many traditional telephone companies used to sell telephone sets. In view of this, why should an Internet-service provider not also sell routers, firewalls for security, etc.? Service providers are now in a position to provide one-stop shopping to a growing and willing market. Other opportunities include server-hosting or Web-hosting services, security services, content-development services, and corporate remote-access services.

The Technology

As shown in *Figure 6*, most access to the Internet today is by conventional facilities: dial-up or dedicated lines, the public switched telephone network (PSTN), or integrated services digital networks (ISDN). Access may also be gained through newer digital subscriber lines (xDSLs) or cable modems. To

TABLE 1

Faster Than Modems

ISDN, ISDL, HDSL, SDHL, ADSL, VDSL

	Data / Voice	Data Only
Shares Voice Pair	ISDN	ADSL
Separate Data Pair	--	IDSL SDSL VDSL

date, these technologies have not produced a great deal of revenue. They are still of interest, however, as they provide higher-speed access. As many people move to small or home offices, the need for high-speed access is increasing (see *Figure 7*). The best future applications will provide higher bandwidth and higher-speed access to the Internet.

Table 1 helps differentiate new access methods in the telephone network: ISDN, asymmetric digital subscriber lines (ADSL), ISDN digital subscriber lines (IDSL), symmetrical digital subscriber lines (SDSL), and VDSL, an updated version of ADSL. (Most recently, a simplified form of ADSL—sometimes called ADSL lite—has been proposed.) For higher-speed access, especially in the small-business and the home-office environments, these are good technologies. They are all faster than a normal modem on the telephone line. The distinction to notice is that ISDN, a switched service, can provide the same functionality as a normal voice telephone line. Hence, if the user has only one pair and wants to use both voice and data, this is an option. If the user wanted to use a separate data pair so that there was one pair of wire coming into the house for the telephone and another pair for data purposes, then there are several different choices. Another important distinction for ISDN is its ability to mix data and voice. ISDN is the only higher-speed service that offers this capability. With ISDN, it is possible to use both of the B channels for telephone calls at one time and bond them together to have a 128-kbps data connection at another.

FIGURE 7

Focusing on the Home? High-Speed Access

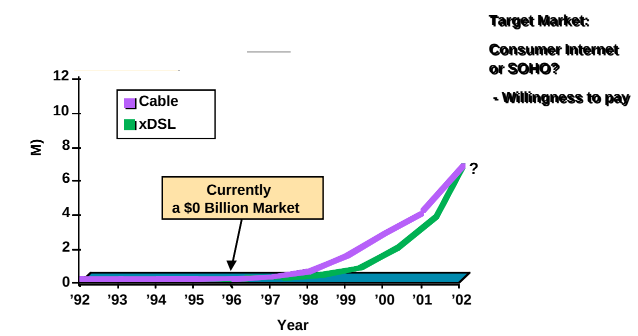

Source: Internal estimates based on current Internet growth rates and customer plans.

FIGURE 8

ADSL

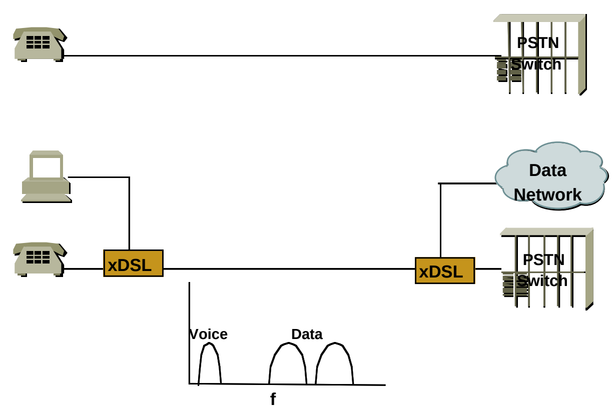

FIGURE 9

High-Speed Access Is Here

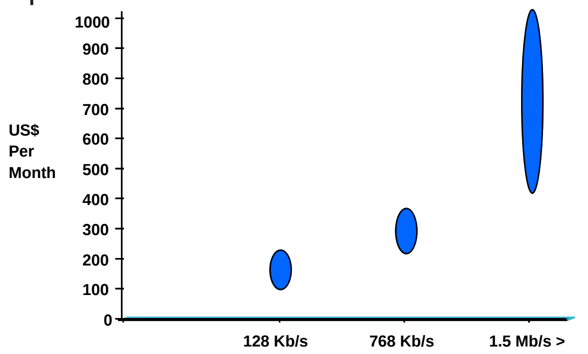

Sources: US West, WorldCom, Warp Wireless Access, HarvardNet, Signet (Texas), Network Access Solutions, ATU Telecom (Alaska), Internet Global Services (Texas)

In the data-only category ADSL is the only technology that uses the same physical pair as the telephone, although it is logically separate (see *Figure 8*). All of the others—IDSL, SDSL, and VDSL—require a separate pair, and sometimes two additional pairs. Regardless of these differences, there is plenty of available technology to meet the increasing needs of the Internet. Much of this technology is already available or in trial (see *Figure 9*). The challenge lies in using it to make a business of the Internet.

Internet Applications

As discussed earlier, it is wise for the service provider to look beyond simply providing access to the Internet, because access to the Internet is well on the way to becoming a commodity in itself. It is not a very impressive or differentiated service. *Figure 10* lists Internet-interface provision near the bottom of the hierarchy of services. It is easy for any service provider to provision base facilities or leased lines, whether in the loop plant or in the interexchange plant.

What becomes important is differentiation. While lower cost may work for a time, it does not compare to providing more services. Services, rather than commodities, offer more opportunities for differentiation. In addition to transport services, companies should consider providing internetworking services and application services. With regard to technology,

FIGURE 10

Service Model

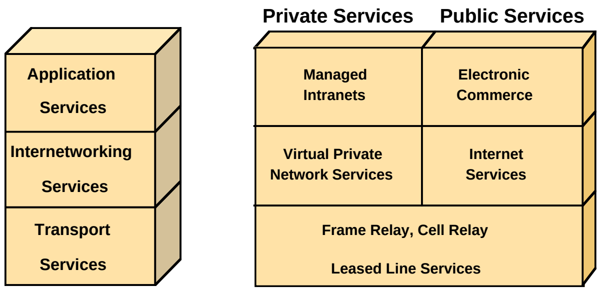

FIGURE 11

Moving Up the Value Chain

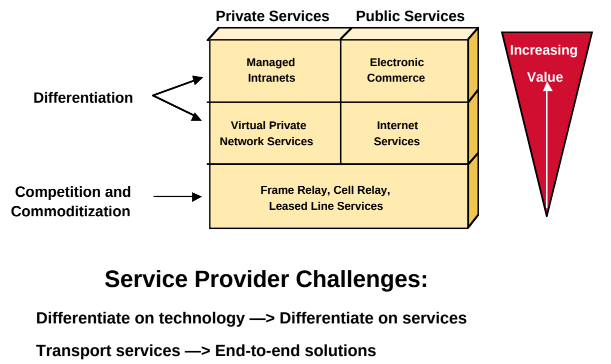

they can offer frame-relay, cell-relay, and leased-line services. Beyond the technology, however, it may be lucrative to cater to the private sector, offering virtual private network (VPN) services and managed intranets, and to the public sector through Internet services and electronic commerce.

Figure 11 shows how it is possible to move up the value chain from areas of competition and commoditization to areas of differentiation. Service providers should challenge themselves to differentiate through services as opposed to technology and through end-to-end solutions as opposed to transport services. It should not be difficult to offer tiered service levels on the same network, multimedia support, and a variety of billing options. By providing premium delivery of content, a business-class-Internet-provider service should be easily differentiated from the standard Internet service.

One way to differentiate is through the intelligent Internet, (see *Figure 12*). In the same way that the intelligent network (or the advanced intelligent network) has helped the telecommunications provider, the intelligent Internet will allow them to offer very attractive services, such as service definition

FIGURE 12

Differentiating the Services

Intelligent Internet

- Service definition unique to each user
- Service creation by service provider
- Dynamic control of network resources
- Consistent across multiple access technologies

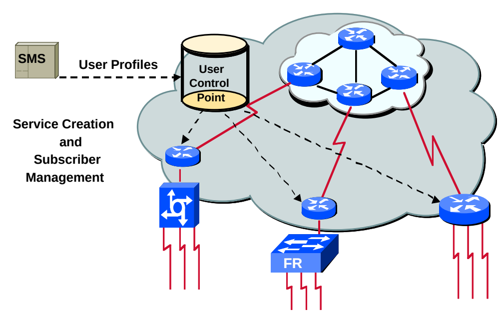

FIGURE 13

Lowering the Barriers to Consumptions

Lowering the Barriers to Consumption
CPE Packages

- Compliments Internet or VPN service
- End-to-end solution approach
- Routers, switches, muxes, gateways, firewalls, servers, software
- Co-branding of package and software

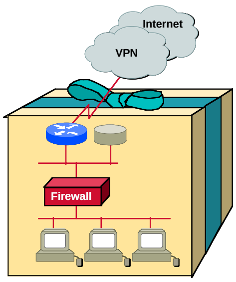

TABLE 3

Web Server Hosting

- **Meets needs of medium and small businesses**
- **Server Complex:**
  - UNIX or NT servers
  - Load distribution
  - Multiple access paths
  - Management
  - HTML and Java tools for development

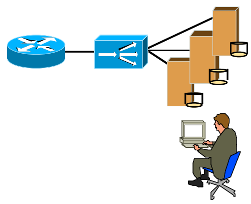

unique to each user; service creation by service provider; dynamic control of network resources; and consistency across multiple access technologies.

Entering New Markets

In addition to offering premium services, providers may wish to consider other areas in which revenue could be gained. One of these is providing customer premise equipment packages (see *Figure 13*). This would be particularly helpful in marketing to smaller companies. For example, instead of offering an Internet access account to a company using 20 or 30 PCs, a package could be accumulated that consists of a small Web server, a security firewall, data routers, switches, multiplexors, gateways, etc. This would offer the company an opportunity to have its presence on the Web for advertising purposes and the ability to deal with the results. Such a package, which would complement Internet or virtual private network service, might even include employee training. By offering this kind of end-to-end solution approach (see *Table 2*), a provider could co-brand software and the package, attracting new and loyal customers. Web hosting (see *Table 3*) is another option and a service that is increasingly popular with business customers.

Another promising way to break into new markets is through Internet telephony. *Figure 14* shows the number of facsimiles sent via long-distance calls versus intercorporate use, and the opportunity for cost savings is great using the Internet for this purpose. This is an example of enhanced services that Internet service providers can offer in addition to e-mail and Web access. Other network-telephony alternatives may become popular in the future, but there are some technical problems to be worked out first.

TABLE 2

Corporate Remote Access

- **Package for corporate location, plus packages for each user**
- **Access servers**
- **Virtual Private Network (Layer 2 tunneling)**
- **Security:**
  - Authentication (users login ID and password)
  - Authorization (users' access privileges)
  - Accounting / auditing
  - Packet filtering (by applications or IP addresses)
  - Encryption

Internet telephony functionality is frequently limited to PC-to-PC-single-point connectivity, resulting in AM-radio quality and poor latency. However, strong economic justification will drive companies to exploit non real-time telephony opportunities such as the facsimile use mentioned above, voice mail, and Internet-to-PSTN gateways. Companies will also be induced to overcome quality barriers in order to provide improved latency and quality. And new carriers, plus many traditional ones, are building dedicated IP networks to offer services with near–PSTN quality.

Security

Security on the Internet is a great concern. Many people have spent a large amount of time finding ways to make the Internet safe for confidential transactions. One promising solution is the use of firewalls (see *Figure 15*). If a service provider considers firewalls, they must understand what type they should offer. It is unnecessary to spend a great deal of time and money creating a firewall for a user with little to protect. Fortunately, the technology is readily available to provide any level of security. This should enable providers to create custom packages for each corporate location plus packages for each user. Security measures include authentication, authorization, accounting and auditing, packet filtering, and encryption.

Conclusion

The Internet is sweeping the world and shows no signs of slowing down. The choice that providers have is to be overtaken and overwhelmed by this trend or to use it to their best advantage. It is important to provide access to the Internet, but if this is where a company stops, then it is limiting itself to the commodity level. There are over 4,500 Internet service providers in the United States—it is only through adding enhanced services that providers will differentiate themselves.

FIGURE 14

Breaking Out into New Markets

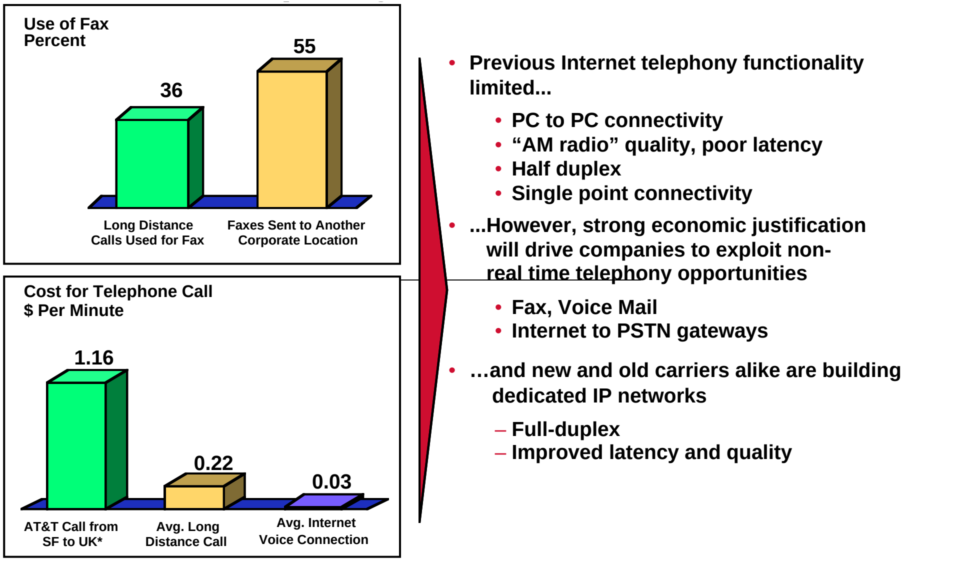

*5 Minute Call at 11a.m., Basic Corporate Rate
Source: ACTA; ATT customer service; Gallup Survey of Fortune 500 fax use.

Providers should also capitalize on ubiquitous presence and service reputations. They should provide the service and comfort that small- and medium-sized businesses—the largest potential market—would welcome.

It is also important that service providers do not neglect their own intranets (see *Figure 16*). Telephone companies have opportunities to deploy intranet capabilities and make a significant amount of money in operational savings. This is true primarily because of the Internet. For many years, technicians have been focused on having all computers interact, but they have had little success. The Internet has solved this problem—any computer operating on any program can access a Web page coming from a completely different operating system. Through the use of Web browsers, companies can save a significant amount of money in training on new programs as well.

The Internet promises to continue to provide a wealth of opportunity. The most important matter for hopeful providers to remember is that they should always be thinking beyond the obvious and the simple, trying to gain and maintain market share through enhanced, customized offerings.

FIGURE 15

Security Internet Firewalls

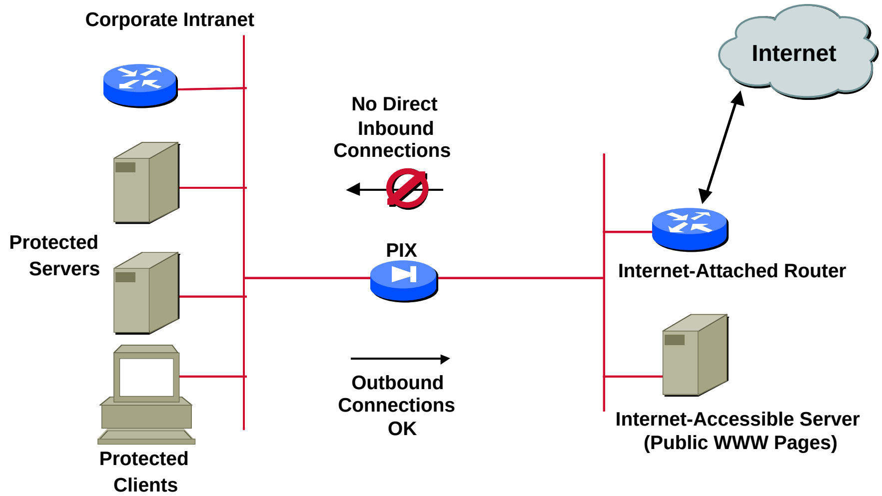

FIGURE 16

The Intranet Explosion

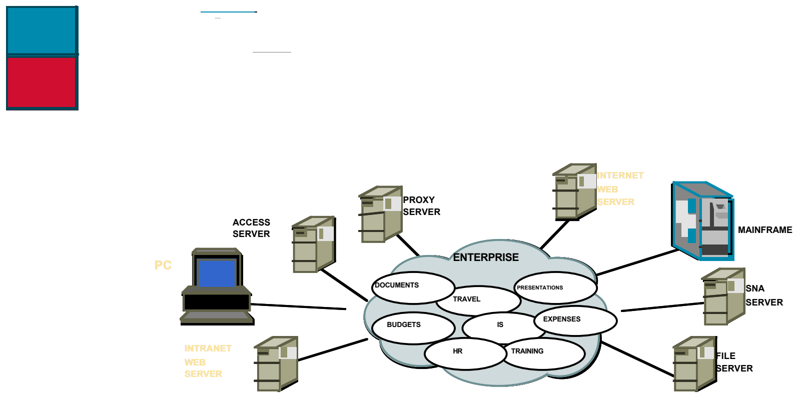

Universal access to all corporate data via ubiquitous interface...the Web Browser

The Accidental Standards

Clients		*Servers*
Windows	HTML	Unix
OS/2	HTTP	NT
Macintosh	TCP / IP	MVS
Unix		VMS

Industry Life Cycles: A Simple Framework for Understanding the Internet

Ronald Hubert
Partner
Deloitte & Touche
Consulting Group

Louis A. Rodriguez
Senior Manager
Deloitte & Touche
Consulting Group

Industry Life Cycles : A Simple Framework for Understanding Complex Markets

The choices and challenges for business managers and investors are becoming more complex. New competitors, technologies, products, and organizational relationships, as well as the pace of innovations are driving fundamental changes in our business environment. The Internet and other new media markets, in particular, are viewed as volatile, complex, unpredictable, and representative of new models of competition.

This paper will describe how these models of competition are not really new at all. The nature of competition and the "rules-of-the-game" have remained relatively constant, in spite of the dramatic changes in the "speed-of-play" and other influences affecting the business environment. Industries and markets follow recurring and stable developmental patterns that are, in fact, more "predictable" than they seem. These recurring "life cycle" patterns can be applied to the Internet and other apparently complex markets. They can be used to describe, explain, and predict the nature of competition, as well as the key success factors and dynamics of winning and losing in these markets.

An industry life cycle is a conceptual model used to understand the evolution of markets, the stages of development, and the competitive dynamics inherent at each stage. While certain attributes of the Internet marketplace are new and revolutionary, the underlying dynamics of market development, evolution, and competition are old and relatively stable. Markets emerge; they grow, mature, and decline. Each of these life stages has different demand factors, levels of sales and profits, and customer dynamics. There are also differences in the factors critical to success in each stage of the cycle; in the competitive strategies that succeed; and in the infrastructure platforms required. Frameworks such as product, demand, and technology life cycles are related and describe different applications of market evolution.[1]

The four basic stages of the industry life cycle are (a) introduction, (b) growth, (c) maturity, and (d) declinee. A general industry life cycle is depicted in *Figure 1* and is similar in overall structure to product and technology life cycles.

For each stage of market evolution, a series of strategic questions may be asked:

- How do we characterize the market? What are the underlying market dynamics?
- What is the strategic objective of the market participants?
- What is the nature of competition? How many competitors will prevail?
- What are the key success factors? How should market participants compete in terms of products? scale of operations? organization? degree of vertical integration? channels of distribution?
- What marketing strategies are appropriate in terms of pricing? packaging? promotion? mass-market vs. niche strategies?
- How will competitive initiatives impact market evolution, and how can a company position for future success?

The answers to each of these questions will vary based on the stage of market evolution. Effective strategies during "introduction" will be different than what works during the "growth" stage. The different characteristics of each stage of an industry's life cycle are described below:

The Stages in an Industry's Life Cycle

Introduction Stage: Emergence of a New Industry. New industries can emerge in a variety of ways. They can result from breakthrough new products or services intended to meet some suspected but latent consumer demand, such as Federal Express overnight package delivery. Alternatively, a new industry can emerge from technology breakthroughs that create a new market or redefine the fundamental economics of existing ones. Radio, television, the telephone, the computer, and many other basic innovations of the industrial age are examples of this source of new industries.

FIGURE 1

Stages of an Industry Life Cycle

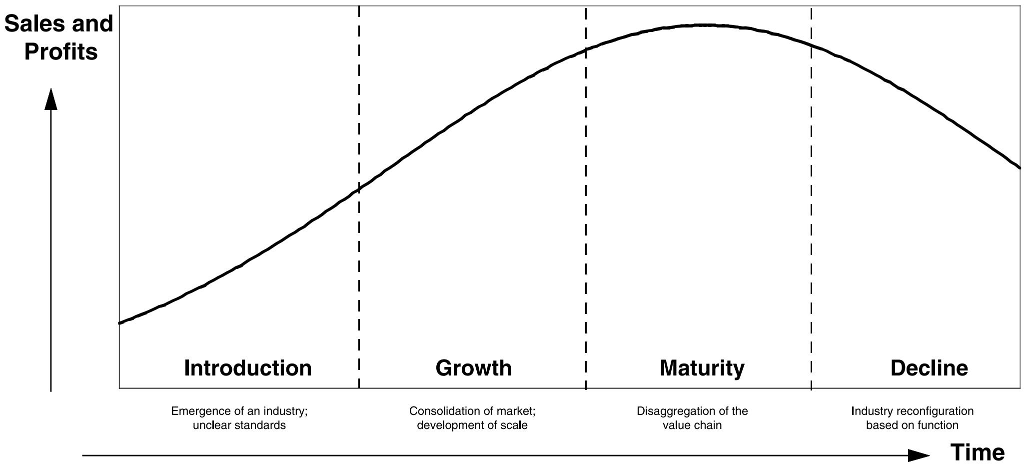

Source: Deloitte & Touche Consulting Group, 1997

In the initial stages of a new market's evolution, companies compete for "early adopters," often with substantially different products (i.e., different product designs, underlying technologies, proprietary advantages, and distinct patents). The "introductory" life stage is characterized by unclear standards. It involves testing and trial of alternative product designs and market approaches with a view toward positioning to control market standards. It is likely that there will be a number of competitors, each with a distinct value proposition and an attempt at product differentiation. Until a "standard" emerges, market demand is relatively low. Subsequently, return on investment is low as market participants balance product development costs, market entry costs, and low-scale production against the limited market revenue available.

Growth Stage: Development of an Industry/Consolidation of the Market. Once a standard begins to emerge, the nature of competition changes and the market evolves to a new stage. In this stage, the requirement is to consolidate a leading market position around the standard established. This requires focusing on growing the market, penetrating the potential buyers, and capturing a dominant share of the buyers attracted. Consumer take-up can occur more rapidly than in the "introduction" stage because the product becomes an established solution for growing needs. This requires that "standards" are more clearly defined, the product or service is well-promoted, the pricing is adapted for large or mass-market acceptance, and distribution channels are arranged to provide for an easy purchase. Infrastructure must also be created to support large volumes of consumers buying the product. Subsequently, the market grows.

The competitive dynamics required to achieve this growth are challenging. A competitor must believe they have defined the clear market solution or technical standard: the market "standard." The competitor must then invest—and risk—substantial capital to build the infrastructure to support rapid mass-market growth that will, in turn, bury the competition. Depending on the speed of adoption, this stage of market evolution tends to be dominated by just a few players who invest to create scale of operations, vertical integration, and wide channels of distribution.

There are a number of historical examples of market consolidation. John D. Rockefeller moved horizontally and vertically to consolidate the oil industry around Standard Oil in the late 1870s. He built huge scale in operations and developed dramatically lower costs than his competitors. He controlled the value chain, the channels of distribution, and ultimately the market for oil in the United States. Henry Ford also moved horizontally and vertically to consolidate the auto industry in the early 1900s. By 1921, Ford controlled over 60% of the new car market. In recent years, well-established industries such as beverages, software, personal computing, and others have exhibited consolidation around a few major providers.

Maturity Stage: Disaggregation of the Value Chain. As market penetration increases and the market grows, several factors begin creating instability. The continuing dominance by a small group of competitors often provides insufficient alternatives to meet the growing variety of customer needs. At the same time, the market is large enough to support stand-alone niche players that can focus on customer requirements that are not adequately serviced by the dominant players. At high enough volumes, portions of the value chain become large enough processes that market participants may focus on only part of the value chain to provide specialized services. Technologies and sub-technologies can be developed to support specific areas of this value chain. As a result, the value chain begins to fragment as suppliers and competitors focus on specialized areas within the marketplace.

At high enough volumes, portions of the market also become large enough to justify their own product offerings to end users. At a high enough level of viewing, for example, cable networks were able to penetrate the broadcast industry and

FIGURE 2

Industry Lifecycle Characteristics

	Introduction	Growth	Maturity	Decline
Market Factors				
Market Description	• Emergence of an industry - Unclear standards	• Consolidation of market; development of scale	• Disaggregation of the value chain	• Industry reconfiguration based on function
Nature of Demand	• Low sales, volumes, profits	• Increasing sales and volumes	• High sales / volumes; competition for profits	• Flat or decling sales, volumes, and profits
Core Customers	• Early Adopters	• Mass-Market	• Niche Consumer within the Mass Market	• Late Adopters • The Commodity Buyer
Number of Competitors	• Many	• Fewer major players	• Many	• Fewer major players
Basis of Competition	• Product differentiation	• Mass-market product • Mass-market scale	• Specialization • Product portfolios	• Commoditization
Market Strategies				
Strategic Objectives	• Establish market standards	• Drive volumes • Dominate share	• Own high-value markets • Multiple niche strategies	• Reposition products • Harvest
Key Success Factors	• Distinct value proposition • Proprietary advantages in product / service • Technological superiority • Patents	• Infrastructure to support mass-market volumes • Large scale operations • Vertical integration • Wide distribution channels	• Infrastructure to support niche markets • Focused capabilities • Specialized expertise	• Customer/product profitability • Operational efficiency • Evolve new technologies and products
Market Positioning	• Drive product awareness	• Mass-market pricing • Market penetration	• Niche marketing/micro-marketing	• Low prices

Source: Deloitte & Touche Consulting Group, 1997

achieve a satisfactory return on investment. Niche products, such as cable networks and specialty magazines, evolve as new competitors begin to carve out increasing shares of the market. In this environment, the number of competitors is likely to grow and become more highly specialized.

In the face of focused competitors, vertically integrated companies can be at a disadvantage trying to provide a full range of products to all customer segments over the entire market geography. Once a market is developed to support focused competitors, these competitors can take advantage of the fact that different portions of the value chain often require different business models. The critical success factors for managing capital-intensive production facilities or commodity product infrastructures differ significantly from the drivers of success in marketing, customer service, and product development.

In a recent article, the *Wall Street Journal* called these focused competitors "category killers," describing them as "Relentlessly specialized and ruthlessly efficient, 'category killers' destroy or devour rivals who lack their single-minded focus." Under these conditions, the value chain becomes more fragmented, and the industry trends toward greater disaggregation along elements of the value chain focused on specific market needs.

Decline Stage: Reconfiguring the Industry. At this point in market evolution, industries generally evolve in one of two directions. If the creative turmoil of a market's fragmentation creates technology or service breakthroughs, the market shifts to a new stage of emergence, struggling to establish new standards and business models for the redefined environment. On the other hand, markets not recast into a new introduction phase will tend to reconsolidate, this time along functional lines. Infrastructure service providers or capital-intensive production facilities where economies of scale dominate will normally merge, as will marketing and customer service organizations. A summary overview of the key strategic questions and the respective stages of market evolution are provided in *Figure 2*.

The Internet Marketplace: Challenges for Success

A number of large, well-funded, highly sophisticated competitors have struggled in their Internet strategies. GE, IBM, Time Warner, CompuServe, Prodigy, and other major players have failed to find the key to success in this marketplace. Many of the initial market entrants, Netcom, PSINet, and others are still looking for a long-term successful business model. Other independent Internet-based organizations, such as Amazon.com—an Internet book retailer—have found that alliances are required to sustain their business model.

What is the economic model for the Internet? What are the important success factors in the Internet marketplace? The conventional view of the Internet has been clouded by substantial media hype, myths, legends, and a new and somewhat confusing set of technologies, products, and market relationships. The "speed-of-play" has contributed to market confusion as apparent paradigm shifts seem to take hold every 30-45 days, rather than over 30-45 years. As such, the Internet and other new media markets, are viewed as volatile, complex, and unpredictable.

Nevertheless, managers can break through the "noise" attached to the Internet and other new markets and use simple, commonsense frameworks for understanding competitive actions. These frameworks include the conceptual model of industry life cycles described earlier.

An Industry Life Cycle View of the Internet Marketplace

Introduction of the Internet: Emergence of a New Industry. In the "introductory" stage of an industry, one would expect to see a number of competitors offering differentiated products to a relatively small group of "early adopters." These distinct value propositions would be characterized by different product designs, underlying technologies, and proprietary advantages. This would result in a lack of clarity around the market "standards" from both a technological and business model perspective. In the "introductory" stage of an industry, one would also expect market participants to be judicious regarding their operations and infrastructure, given low volumes.

The Internet marketplace emerged along these lines. The "Internet" itself evolved from academic and military research in the late 1960s. Through the early 1980s, the Internet connected a small community of researchers. At this time, a diverse stew of organizations had begun to provide additional electronic information services. For the most part, these bulletin boards, discussion groups, electronic libraries, and on-line services, including AOL, CompuServe, Prodigy, Genie, Delphi, provided distinct offerings within closed environments. Each had a different underlying economic model.

This evolution is similar to other industries. Initially the pursuit of hobbyists, the personal computer industry emerged from a number of competing alternative technologies including CP/M, video game-related systems from Commodore, Atari, and Tandy, and the initial landmark work from Apple. All of these options were complex to set up and use, and almost all were incompatible with each other.

Windows 3.0 for IBM-compatible personal computers was launched in 1990, establishing the interface which would become the "standard" for personal computers, on-line services, and, ultimately, the Internet. The World Wide Web, a Windows-based point-and-click mechanism for navigating the Internet, was released in 1992. A number of academic and entrepreneurial organizations developed additional Windows-based Internet tools—browsers and search engines—to support navigation on the Web. Subsequently, government, the entertainment community, and the business community at large began to place large amounts of information on the Web, leveraging the new Windows-like mechanisms and linkage tools. Windows became the "standard" interface for the Internet.

AOL became the leading on-line service provider because it adopted a Windows-based interface sooner and more effectively than its competitors. It began to move toward a mass-market model of competition, with mass-market advertising, and an open environment that allowed for Internet browsing. Although the market leader, CompuServe, had established a strong position among "early adopters," including researchers, academics, and technophiles, the basis of competition had begun to shift toward mass market, and CompuServe did not adapt fast enough. The stage was set for market growth. This growth was enabled by the Windows "standard." The early adoption of this standard by AOL positioned it for growth.

As the growth stage began to emerge, market participants began to position their product offerings, capabilities, and infrastructure further along the value chain. Many of the early on-line service providers who had leased facilities and infrastructure for data transport, had an eye toward vertical integration, while backbone providers and infrastructure stalwarts (i.e., MCI, Sprint, as well as UUNet, BBN, and others) evolved retail-based solutions, particularly for the business market. The stage was set for market growth, vertical integration, and scale.

Growth of the Internet: Development of an Industry/Consolidation of the Market. In the "growth" stage of an industry, one would expect to see offerings intended for the mass market in terms of pricing, product, and packaging. These products would be offered by larger, more vertically integrated players with the infrastructure to drive mass-market growth. They would need high-capacity distribution channels to achieve large volumes quickly.

In spite of its relatively small size, AOL positioned itself to drive market growth. From 1993–1995, AOL grew from just over 300,000 subscribers to 4.5 million, taking over leadership in the on-line services marketplace. AOL expansion in subscribers had allowed it to increase revenue from $50 million in 1993 to $400 million by 1995. Its scale of operations and marketing infrastructure had increased substantially. AOL created vertical integration to achieve market dominance through its acquisition of ANS, one of the original backbone infrastructure providers, and its negotiation of significant transport relationships with BBN and Sprint. AOL also began to solidify relationships with key content providers as the advantages of mass-market volumes became clear to the production community.

AOL adopted mass-market pricing, following the lead set by AT&T. Leveraging its 70 million household customer base, AT&T launched a pricing strategy designed for mass-market adoption. Unlimited usage had been deployed by several independent ISPs, including Netcom, but not by a major conglomerate. The AT&T initiative established the market price that AOL met and extended with an unprecedented marketing barrage. Over 150 million AOL starter kit disks were mailed during 1996. AOL disks surfaced in mailboxes, in magazines, on airlines, and in promotions of every shape and size. AOL subscriber volumes grew to just fewer than 8 million by 1996 representing 45 percent of Internet subscribers.

FIGURE 3

Value Creation Cycle for Growth Phase of Development

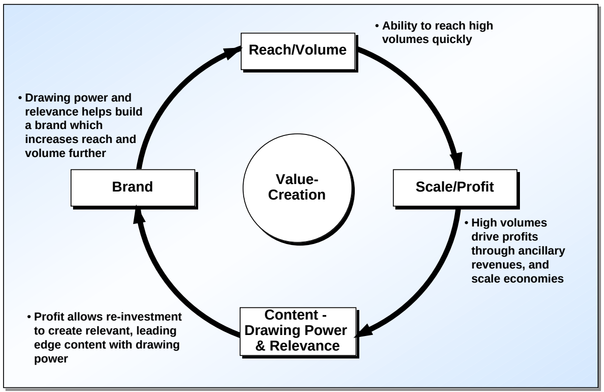

Source: Deloitte & Touche Consulting Group, 1997

AOL established a business model based on scale and market dominance. AOL leveraged a massive marketing barrage to achieve mass-market volumes, achieving 10 million by 1997. It has subsequently generated over $1 billion in annual subscriber fees. At these high volumes, AOL has secured highly lucrative marketing alliances, substantial advertising revenues, and additional premium revenues. This economic model represents a proprietary advantage that is self-reinforcing. Large revenues and profits allow AOL to further develop content relationships and product, build brand and marketing initiatives, and increase volumes that, subsequently, increase subscribers and revenues even further. *Figure 3* details the dynamics of a value creation cycle that supports the growth phase of development.

This evolution is similar to other industries, where the self-reinforcing benefits of scale and market dominance supported development of the oil, auto, telegraphy, and telephony industries. In the broadcasting industry, the foundation for dramatic growth was cemented by large, well-funded, vertically integrated providers with strong distribution channels and marketing capabilities: RCA, General Electric, and Westinghouse.

In 1926, the National Broadcasting Company (NBC) was formed under the ownership of RCA (50 percent), General Electric (30 percent), and Westinghouse (20 percent).[2] RCA, General Electric, and Westinghouse leveraged their relationships with major retailers to sell millions of radios. They developed content, information, and entertainment that they leveraged across a national network. They leveraged the costs of marketing and advertising across radio affiliates throughout the country, and, in turn, generated advertising revenues from marketers in return for mass-market audiences. *Figure 4* details the market penetrations and growth phases of a number of communications-related industries.

Other providers with the ability to achieve mass-market volumes have achieved success during the growth stage of Internet evolution. Both AT&T and Microsoft have leveraged their respective large installed customer bases to build volumes and achieve some of the benefits of scale. Among backbone providers, UUNet has achieved similar market dominance by establishing key relationships across the value chain and leveraging its scale into further consolidation. Its strategic relationships with Microsoft, Earthlink, and other retail Internet providers provides UUNet with large volumes. Its recent consolidation of the AOL and CompuServe infrastructure increases its scale in the industry.

A number of large, well-funded, highly sophisticated competitors have fared less well. The IBM electronic shopping mall shut down. Pathfinder, the content offering from Time Warner has generated substantial costs with limited returns. CompuServe and Prodigy have stalled.

These players did not achieve the infrastructure required to support mass-market volumes, including large-scale operations, vertical integration, wide-distribution channels, and the immense marketing effort required to achieve any substantive share of a growing, yet relatively small customer base of 20 million subscribers.

What's Ahead? Continued Growth? Disaggregation? Reconfiguration? In the "growth" stage of an industry, one would ultimately expect to see mass-market volumes, high penetration, large-scale competitors, vertical integration, economies of scale, operating efficiencies, and high-capacity distribution capabilities. As these players drive toward high volumes, portions of the market become large enough to justify their own product offerings. In the "mature" stage of an industry, one would expect to see focused competitive offerings meeting the demand for niche product offerings. At high enough volume levels, one would expect to see specialization and disaggregation of the value chain.

The Internet marketplace grew 60 percent from 1995 to 1996, growing from 11 million subscribers to 18 million. By year-end 1997, 22 million subscribers had been on the Internet. While this represents substantial growth—double the market in two years—the growth has substantially slowed in 1997 over 1996. Most forecasts look to roughly 40 million subscribers by the year 2000.

By 1997, the average American will watch over 30 hours of television each week. The Internet represents a fraction of that usage. Much larger volumes and higher penetration levels will be required for content providers and the entertainment community, who are currently receiving little or no income, to achieve any substantial returns.

The Internet marketplace is still in its growth stage. If the market continues to grow, the consumer should expect to see more consolidation, more vertical integration, and the continued development of the operating infrastructure to drive much higher volumes.

There are over 4,000 ISPs in the United States today. These largely independent entities represent an additional distribution channel for increasing the Internet subscribers' base and a congestion bottleneck in the Internet infrastructure and a limitation to mass-market growth. The expected consolidation of this market may support the development of the market.

FIGURE 4

Market Penetrations of Communications Related Industries

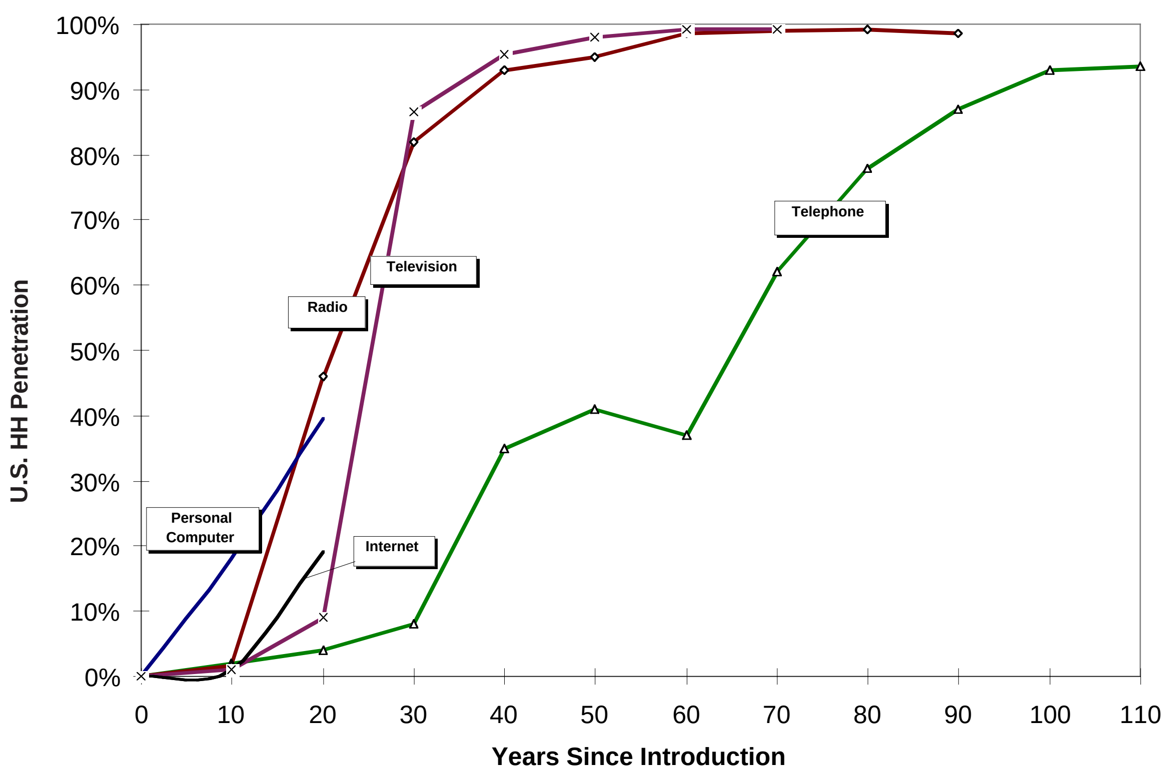

Source: Insight Research Corporation, 1996; Veronis Suhler Associates, 1997

The enormous infrastructure investment and consolidation required to develop next-generation broadband transport technologies which will carry both voice and data traffic may represent the final operational requirements to drive high-volume growth in this new marketplace. It could potentially supplant the traditional telephony-based industry model and drive Internet-based applications beyond 40 million subscribers.

The early market signals from AOL, the sale of ANS, and the new alliance with UUNet may signal an early view of market maturity and disaggregation of the value chain. On the other hand, it may signal a cohesive alliance between players with vested interests, including Worldcom, MFS, UUNet (also the backbone provider for MSN) and AOL and Microsoft. Together, these large-scale players represent much of the future of broadband applications.

Conclusions

What is the economic model for the Internet? What are the success factors in the Internet marketplace? The answer is: It depends on the stage of industry evolution. For the early years, product differentiation and positioning for market standards will be the keys to success. In the "growth" stage, large-scale competitors, vertical integration, economies of scale, operating efficiencies, and high-capacity distribution capabilities are imperative for success. The current winners have operated along these lines.

In the "mature" stage of market development, one would expect "category killers," who are relentlessly specialized and ruthlessly efficient, to destroy or devour rivals who lack focus. In the "decline" stage, market winners position to reposition and reconfigure their business model to evolve into a new industry model.

References

1 Perspectives on "Industry Life Cycles" are based on concepts regarding market evolution, demand / technology life cycles, and product life cycles from Kotler, P., *Marketing Management: Analysis, Planning, Implementation and Control.* Englewood Cliffs, NJ: Prentice-Hall, 1991, Chapter 5, pp. 348-372.

2 Paper, L., *Empire: William S. Paley and the Making of CBS.* New York, NY: St. Martin's Press, 1987, p. 19.

Internet Over Broadband

Jim Kaubisch
Media Architect, Network Products Group
Sun Microsystems

For Internet service providers (ISP), the Internet is a maze of opportunity. Everyone is racing for new technology and new services, as every step along the technological road is potentially a new revenue source. In order to make money, an ISP must supply or deploy responsive and reliable services and service infrastructures that result in user satisfaction. If the quality of service is not there, the users will leave and the ISP will not make money; if the system is not properly managed, the ISP does not have a good idea where to put resources and / or how to deploy them to maximize customer satisfaction or provider profit. The purpose of this paper is to illustrate some of the challenges that face operators of broadband networks—such as cable television operators—planning to deploy Internet services over their network and to give an overview of a possible infrastructure that could meet those challenges.

Considerable complexity underlies this overview, of course, and throughout this paper it will be evident that no one truly understands where Internet services over broadband are heading, what all the business and technical issues will be, nor that the approach outlined here or any alternative approach will completely meet the needs of the future. However, the hope is that this discussion does present a degree of insight which will help in understanding the debates currently being waged. Before addressing the issues, however, it is perhaps useful to have a look at what Internet services comprise and how they might evolve.

Services and Markets

The Internet is not something that happened three years ago; it has been around for approximately 30 years. The Internet was first set up by the United States Government's Advanced Research Project Association (ARPA) as a noncommercial network to share network resources (then called ARPANET). In the early 1980s, transmission control protocol / Internet protocol (TCP / IP) was developed and deployed as the protocol of ARPANET and its inclusion in the University of California at Berkeley's Berkeley Software Distribution (BSD) version of UNIX meant that the number of machines which had access to the network greatly increased; this newly expanded interconnection of networks, based on TCP / IP, was renamed Internet—it is the interconnection network.

For many years, the number of people using the Internet could be counted in the thousands. Why did it suddenly gain general popularity and increase its user base to many millions? A so-called "killer app" was developed called Mosaic. A "killer app" is an application which in some way causes a revolution in the way people think or do business. (What would the course of the IBM PC have been if Lotus 123 had not been invented?) Using the hypertext markup language (HTML) to devise information pages which could have dynamic links to other pages, and the hypertext transfer protocol (HTTP) to request and transfer the pages, Mosaic was the first application to provide the graphical user interface (GUI)–based access to the information on the Internet that everyone is now familiar with through other "browsers" such as Netscape Navigator and Microsoft Explorer. Suddenly, the wealth of information already available on the Internet infrastructure and new content developed for access by the GUI browsers became available to the common person. Before Mosaic was introduced, the Internet was opaque (and unknown) except to those few who knew the detailed hieroglyphics required to gain access. In the few years since Mosaic was invented, the whole nature of the Internet has changed and continues to change at a dizzying rate.

The Internet offers many data services: e-mail, news groups, Web access, Web publishing, chat groups, work-at-home and learn-at-home potential, audio and video content, interactive games, electronic newspapers and magazines, financial services, and others. Although these services are very familiar, an important point should be stressed: despite the great interest in Web access, other non-Web applications such as work-at-home are likely to surpass Web access in terms of revenue. Home consumers decide what Web access is worth to them. Work-at-home, on the other hand, has a component in which the employer has an interest. The ability to work from home is of value not only to the employee, but arguably of even more value to the employer. It has a measurable economic payback. For example, some states (including California) have proposed or implemented laws which mandate that, over a certain period of time, a percentage of the work force must work from home because the number of parking spaces allowed for a particular building will be reduced in order to decrease the environmental effects of commuting. In many cases, people's jobs require the use of (networked) computers and most of those will have tools with graphical interfaces which, in turn, means that there is a need for high-bandwidth connections which can be most effectively supplied by broadband.

Distance learning (whether at work or at home) is another important market. It has been said in the Silicon Valley computer industry, for example, that the skills of 20 percent of employees become out-of-date every year. If nobody did any training, the entire work force would have outdated skills at the end of five years. The same is likely to be true for many

FIGURE 1

SBIS Architecture

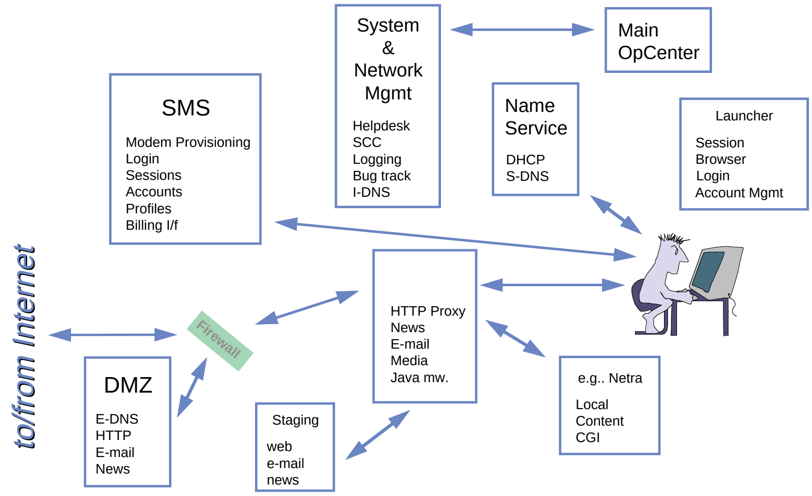

other industries. Technology and competition in markets are moving ahead at such a quick rate that continuous retraining is an absolute requirement.

Retraining has traditionally meant assembling a group of people in a classroom setting for a number of days. Taking employees out of the workplace in this way, however, is a very expensive way of training. The cost of wages while the person is not productive, the lost production itself, the cost of the training facility and instructors, and the cost of travel and accommodation add up to substantial amounts of money and are driving corporate training organizations to look for alternatives. Learning over broadband can meet or exceed the effectiveness of the classroom situation. Broadband can deliver video, voice, or animation to demonstrate concepts. In fact, broadband can present a teacher in a room remote from the student so the learning experience can be more like a standard classroom situation and less like a slide show, yet can be done at times and at a pace which better suits the individual employee.

A development which has garnered a great deal of attention in the past few months is "push" services. Traditionally (if something that has been around for only two or three years can be considered to have a tradition), if you want information from the Web you point your Web browser at the universal resource locator (URL) of a desired Web site and "pull" the information down from the site. More recently, services such as Pointcast have appeared which allow users to specify topics of interest to them when they subscribe and it will "push" the updated information to them which the service itself has pulled from the relevant sites. What the user ends up with is a sort of *Reader's Digest* version of the various topics, often with links to the original sites so they can "pull" the complete story if they desire. However, if their Internet access is via a dial-up line, they can only update information whenever they connect to the Internet, and it can take considerable time to get all the updates at 28.8 kbps. On broadband, where users can contemplate staying connected indefinitely, they may receive updates continuously—a real benefit for watching the stock ticker or seeing the late-breaking industry news which affects their businesses.

Potential new killer apps are emerging every day. The latest one may be Internet telephony, often called I-phone or Web phone. With I-phone, telephony does not require a telephone switch or other device at the head end; it simply requires the two parties who want to have a conversation to have (compatible) I-phone software on their machines. It is true that both must have their computers switched on and connected to the Internet at the same time. There are potential latency issues in this scenario that one would expect in a packet-switched environment similar to the Internet, but as these issues are worked out the I-phone may prove to be a killer app. Would anyone doubt the attraction of being able to make a long-distance call at no cost? For better, more reliable service people are willing to pay a charge. To use I-phone a user needs to access the Internet. When is the last time you got a busy signal when that user dialed his ISP? His answer could well be hours ago and is probably not more than days ago. With a cable modem, users can contemplate being connected all the time and therefore initiate a call at a moment's notice as well as being able to receive a call at any time.

Companies are developing "on / off ramps" for I-phone. Users initiate their calls over the Internet, and, at the destination city, the calls are routed via an off-ramp to the regular phone lines and the phone rings as it normally does—long distance at the cost of a local call. This is a business that a cable operator could well be interested in pursuing.

Someone in his or her bedroom or garage will put a Web server onto the Net for the first time tonight, and although there is a 99.99 percent probability that what they provide will lead to nothing, there is also the near certainty that somewhere, sometime someone will do it and it will be the next killer app.

End Users

One should not lose sight of the fact that whatever services might be available, their success will depend on acceptance by the "end user." Different people use the Internet in different ways. Seventy-five percent of my on-line time is spent finding product information. If I hear that a company is doing something interesting or I see a billboard advertising a product which sounds useful, I go to their Web site and see what the more complete story is. I tend to use product support provided over the Internet, especially obtaining software upgrades—a process much easier than calling an 800 number and waiting for 20 minutes. Other people are more interested in news. They can access foreign newspapers and publications or get several versions of a story with international overtones from the press in each country involved. Others buy Web access because they are interested in joining chat rooms, whether to discuss topics of particular interest or just to "meet" other people. Whatever services are provided, we must not lose sight of the fact that they must satisfy a need, whether real or perceived. From that point of view, Internet services are no different than any other form of service provision.

We tend to think of the end user as the home user, but organizations such as libraries, schools, and hospitals are also end users of broadband Internet services. One reason to deploy cable modems is that a number of personal computers in a school or library setting could potentially share the substantial bandwidth the modem provides. People are not working intensively at those machines for eight hours a day, but the institution nevertheless needs to provide simultaneous network access for a number of people. Small businesses and home businesses are important end users in that the Internet provides them broader access to information than they could otherwise afford as well as allowing them to publish in a way which provides them with more coverage than they could otherwise dream of having. How often have users seen an ad or even had it suggested to them in a conversation that they visit www.fooba.com? When they do, the well laid-out page of a one-person company can present as professional a storefront as that of a thousand-person organization at a cost which is a fraction of the alternatives. These professional end users can more easily measure the economic benefit of their Web access as a result of losing it due to system failures. They are likely to be more demanding but also more stable and willing to pay more for a "professional" level of service.

The Challenges Facing Operators

The first part of the problem is that the revenue models for providing Internet services are untested. No one knows how many customers are coming, what they want if they do come, and how much they are willing to pay for it. Designing open systems which are economically viable for small operations but scale to very large ones is no easy task. Given the discussion so far concerning the evolution of the Internet and the new applications people are finding for it, it is obvious that the Internet access infrastructure provided by operators must remain open and flexible to accommodate these new possibilities. No one can say—as the United States patent officer did in 1899—that the government should close the patent office because everything that can be invented already has been invented. That is simply not true. No one has even the foggiest idea about where the Internet is going, and expecting to be able to provide an infrastructure closed to all but the existing (approved) applications is a formula for disaster—the day after the providers finish building the system, the next I-phone comes along and their users find an ISP through which they can access the exciting new service.

The second issue is that no one has yet built a high-speed computer network that has hundreds of thousands and—ultimately—millions of end users and tens or even hundreds of thousands of service providers. The telephone companies have put together very large networks for plain old telephone service (POTS), and on-line services such as America Online can provide 28.8 kbps access to millions. Yet no one has ever built broadband networks that have the bandwidth capable of carrying audio / video, providing Internet access, accommodating telephone service to—and between—a very large number of people. A multiple system operator (MSO), such as Time Warner or TCI, provides services in multiple cities to almost hundreds of thousands of customers, but these are one-way networks—the head end provides the services and the end users consume them. But no one has yet put together the complete network envisioned here, where the end users both consume and produce content and determine the volume of information traffic across the network via individual requests. Put differently, in the current network if providers broadcast the Super Bowl, it makes no difference to their network load whether one person or one million people are watching; in the network discussed here, they have to be able to accommodate the possibility that many thousands will ask them at the same time what the score is and expect a speedy reply. The question remains untested and unanswered: how do providers ensure the responsiveness, reliability, and security of such a network?

Finally—and this is in some ways the most difficult part of the problem—the Internet is anarchy. It is, perhaps, one of the most benevolent and well-meaning forms of it, but it is anarchy nevertheless. The Internet was developed to be open. It has no underlying security. Anybody can get to anything. Nobody owns it. Nobody controls it (though there are bodies such as the Internet Engineering Task Force which work at putting order into the system). Nobody controls what goes into it, laws governing rights of those who access it and liability of those who provide services over it are in their infancy, and intellectual property issues are unresolved. If users thought that the Internet was the most wonderful thing they had ever seen and wanted to write the Internet a $10 check in appreciation, there would be nobody entitled to receive their check. In that sort of environment, how can anything be secure? These are very difficult problems to solve.

An Outline of a Service Architecture

Proxy Caching

One possible approach a broadband operator can take toward providing Internet service is simply to become the provider of fast "pipes," leaving the provision of other parts of the necessary infrastructure to an independent ISP. There is nothing fundamentally wrong with this and, indeed, it can be argued that this is the most sensible approach in the current frenzy of change, especially considering the capital-intensive nature of doing much more. However, history has shown that there are dangers associated with this approach. Pipes are commodities, their economic value is very price-sensitive (usually means low profit margins), and users have no loyalty to them. Users have rational or emotional ties to the providers of services (i.e., those organizations who provide them with "added value"), not to the wires which deliver them. This is the reason that telcos and MSOs alike are looking to move up higher in the food-chain by providing services identified with them rather than someone else. For the remainder of this discussion, it is assumed that the operator is interested in providing more than the pipe.

A reasonable intermediate step between being a provider of pipes and a full-fledged ISP is to provide proxy caching for your end users. In reading reviews of the performance of cable modems, one often sees remarks to the effect of "The modem gave me an Mbit connection but when I accessed my favorite Web site it wasn't any faster than with my 28.8 kbps dial-up." There is a simple reason why that observation might well have been true: the speed of the connection is determined by the slowest link less all the delays introduced by intermediate nodes in the network. A user might be connected to the Internet via a fast cable modem, but if the Web site is connected via ISDN with ten intermediate machines between them then maximum 128 kbps (and in real life less) is the effective speed of the link. How can the provider of broadband Internet access do anything about that? He can provide a proxy cache placed between the ISP and his users (i.e., he makes a copy of the Web pages requested by his users as they are delivered). The first request for a given page will be as slow as ever, but subsequent requests will be serviced by the local proxy (i.e., the local machine which "stands in" for the real provider of the page) and will therefore be very fast. This would be of no use if requests were truly random (i.e., every request is a "first" request), but studies have shown that there are patterns of access and the perceived improvement in access speed can be dramatic.

Service Hierarchy

In order to be open to the next killer app, operators must provide infrastructures that are open to legitimate flow of information to and from the user. At the same time, they must control access to services in order to realize a return on their commercial value. If an operator supplies on / off ramps to I-phone, for example, he would probably want to charge for that separately from "plain" Internet service. A service hierarchy must then be implemented to define service entitlement. In other words, operators must have some way of defining what is meant by a "service" and controlling access to it. The situation is analogous to the service entitlement systems employed by direct broadcast companies and cable operators to block premium and pay-per-view channels, but the implementation is much more complex.

Security Filters

The operator must provide access and security filters to enforce service entitlement as well as to suppress the transmission of packets which can compromise network operation. It is not good enough to provide access to three services for $19.95 if the user can actually access four additional services. Security filters are a necessity, or the network will not last very long. There are people out there who enjoy breaking a system just for the sake of breaking it. There are also naive users who accidentally break systems (often more effectively than those who try to do it intentionally). The Internet is a network with inband signaling (i.e., the packets which control the network go over the same network as packets which carry user data) and was not designed with any security provisions in mind. There is fundamentally no difference between a packet that carries SNMP network management commands and one that carries Grandma's latest recipe for chocolate cookies. An operator must look at packets and determine the need for that packet to access the network on a case-by-case basis and delete those which are not legitimate.

An analogy and lesson as to the importance of all this can be taken from the telephone network. Until a few years ago, the telephone network used control tones to route calls through the network and transmitted them over the same part of the spectrum used to carry voice (i.e., inband signaling) — some users may remember hearing touch-tone-like noises after they finished dialing long-distance numbers. A number of people found out what these tone sequences were and, besides those who used the famous "blue boxes," some taught themselves to whistle the control tones into their telephone handsets — the telephone switches heard and understood and gave those people the long-distance connections they had requested, all at no charge. Today, the phone network uses out-of-band signaling (i.e., control information goes over a network which is separate from that which carries voice and is not accessible to customers). Unfortunately, this option is not available to the operators of Internet services. Instead, methods have to be devised to detect "whistlers" and delete the information they are putting into the network before it does any harm.

Service Infrastructure Scalability

The operator must also provide service infrastructure scalability for responsiveness and reliability. No one really knows if 20, 20,000, or 200,000 people will try to simultaneously access a particular service, for example. Any infrastructure offered must be scalable to cover any of the likely scenarios that may develop. If you do not allow users access to their favorite service or access is slow / unreliable, they will switch to an ISP who does it faster and more reliably—there are plenty to choose from.

A More Detailed Look

While proxy caching, a service hierarchy, access and security filters, service infrastructure scalability, and comprehensive protection of the service infrastructure are basic components of the response to the problem of developing new revenue sources through the Internet, a great deal of detail can be added to that basic picture. The network infrastructure must be such that it can support a variety of functions. It must support multiple client platforms, including PCS, Macs, worksta-

FIGURE 2
Typical Configuration

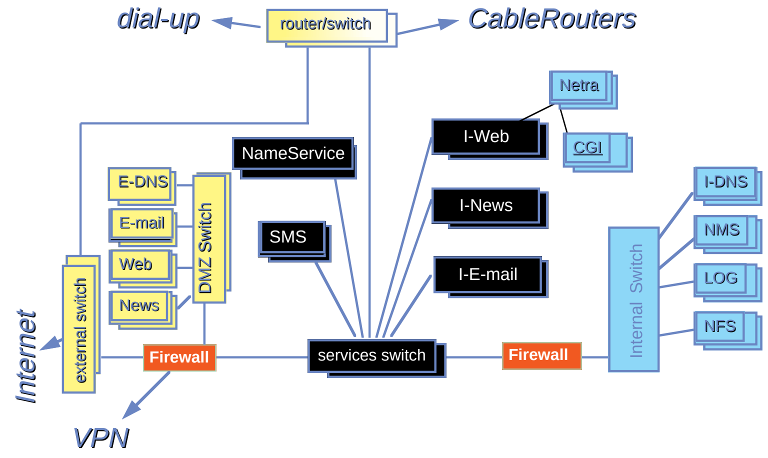

tions, and others yet to be invented. It must provide for server interoperation by supporting multiple service suppliers. It must support all client services, including the Web, news, e-mail, chat, and I-phone. It must provide a wide range of services, including:

- integrated account and log-in session management
- Internet protocol (IP) address allocation and management
- easy administration of the network, systems, and services
- multiple service and "quality of service" tiers
- name services and white pages
- personalized content (per site or per user)
- content creation and management
- dynamic content update (e.g., from databases)

The network infrastructure must also have a variety of important inherent characteristics. It must be secure—the system must protect servers, services, clients, and enterprise links. It must offer high performance, but with control and predictability. In other words, the network must give the performance of a drag racer, the security of an armored truck, and yet have the controllability of a fine sports car. The network must have scalability to accommodate a wide range of user populations, connections, content, processing power, and network bandwidths (if a service does not provide the responsiveness required, the operator can either get a bigger machine or network a number of smaller machines, both of which offer their own advantages). The network must remain available, providing failover to maintain service with the least disruption possible in the event of a system failure. Finally, it must offer reliability through load balancing—to prevent a situation in which one type of service is poorly served while another gets good service—and manageability.

Architecture—Flow of Information

A system architecture like that in *Figure 1* can be developed based on the requirements briefly outlined above. At this level it does not look much different from any other service architecture. The home user wants to get onto the Internet. How is that accomplished? When the users log on, they are connected to the name service where the dynamic host control protocol (DHCP) provides an IP address—this step is in part required because of a shortage of Internet addresses. When the Internet protocol was defined, IP addresses were defined as 32 bits. Based on usage at the time, that was considered enough to cover Internet needs through eternity. Today, however, there is a shortage of IP addresses. The reason is simple: at any given moment an IP address must be unique on the network to which the user is connected, which, in the case of the Internet, means more or less the whole world (analogous to the need for unique telephone numbers). Enough addresses have been given out to produce a genuine shortage of unassigned IP addresses. A development called IP7 is likely to be introduced in the next year or so to expand Internet addresses to 128 bits which will, over time, cause the problem to disappear. As that ability is not available yet, DHCP is a way to temporarily allocate the IP address that an end user needs to access the Internet.

Once a temporary IP address has been assigned, the user is connected to the subscriber management or service management system, which sets up the environment to access what has been paid for: log in, sessions, accounts, and profiles. From there, the user can access various services such as news, e-mail, and similar services—the more dynamic content of the network.

All activities to this point are internal to the network of a particular operator. Getting out to the Internet requires two additional components. First, the operator will install a firewall. The firewall is there to separate the traffic on the Internet from the traffic inside the operator's network. Only certain, specified packets are allowed to penetrate that firewall (this is one of several places where providers try to catch any "whistlers"). The other component is a name service. Access to everything on the Internet requires an IP address; these addresses are tedious to type in and remember and can change when a particular service is moved from one machine to another or is present on a number of alternate machines. To overcome these issues, alphabetic names are usually used such as "www.bestcookies.com" which must be translated into an IP address of the form 255.255.255.255 before they can be used to access the site. If users want to access "bestcookies," for example, odds are that their local machines have no idea how to translate the name into the address. The domain name service (DNS) is a standard way to accomplish the translation. When DNS identifies the IP address of that particular server, the user can establish the connection.

Architecture—Network Components

The simplest form of providing Internet services is to get an Internet server, plug one end into the cable modem system, and the other end into the Internet. This approach offers the minimum level of Internet access and is what many people have. The system configuration shown in *Figure 2*, on the other hand, goes a long way toward meeting the requirements outlined above for an adequate multitiered system architecture. It can be partitioned into four sections:

1. The first section is network operational control to which nothing except network operational traffic has access. There is a firewall here to ensure that no unauthorized traffic reaches this section of the network. A security compromise here could cause a complete network outage as well as destruction of information needed to bring the network back up, meaning that the operator could be out of business for a very long time!

2. The second section is the operator's "private service" infrastructure to which only his subscribers have access. This is the section which would be home to services provided only to the operator's customers such as I-phone on / off ramps. The private service section, by the way, corresponds to what in the Interactive TV discussions was often called "The Mall" (i.e., a collection of services perhaps owned by but certainly operated by the owner of the Mall—for a fee). This section is also home to the name services and subscriber access management components.

3. The third section is the public access infrastructure which is the interface to the Internet as well to virtual private networks (VPN). For the Internet connections, this section has components which are counterparts to many of the components which exist in the private and control sections. There is a DNS server here as well as in the control section, an e-mail server here as well as in the private section, and so on. The reason is that providers do not want to allow direct visibility or access from the Internet into our private network. Instead, providers build a link between the internal and external components with another firewall between them which allows the internal components to accept or reject requests and data from the external (Internet) world. For example, while the external e-mail server may be at considerable risk of being attacked because it is directly visible on the Internet, the internal one is much more protected because it cannot be seen.

   The VPN connection is there to support applications such as work-at-home, inter-office connectivity, vendor connectivity, and so on. As the name implies, the connection from one or more points in the operator's network is connected to a point outside the network in a way that makes the connection look like it is part of a private network—perhaps an extension to a corporate network. These connections are invisible to those who do not belong to them. This is a structure now common in the public telephone network. It used to be that when a corporation wanted an internal wide-area network they would lease lines from telephone operators and manage those networks themselves. More and more, corporations are leasing bandwidth in a shared network in a way which retains the privacy of the line but, in fact, shares the actual physical wires with others.

4. The fourth section is the user infrastructure. This section comprises the components required to connect the end users to the network such as dial-up modems, cable routers, and modems. While not explicitly shown, each of these components has some amount of security filtering to ensure that the user cannot harm the network and, as much as possible, that the network (especially the Internet portion) cannot harm the user. "As much as possible" is used so as not to imply that this security can be made complete. As long as users say that they want direct access to the Internet, the operator cannot provide them complete protection from attacks from that quarter.

With every layer of detail which this paper exposes, more issues are raised and complexities uncovered. Hopefully, though, this discussion has made clearer at a "block diagram" level what needs to go into building one of these systems without leaving the reader with the impression that it is all so hopelessly fraught with difficulty that one should forget the whole thing. Companies are building at least some of these structures, and there are operators who have encouraging results from their first attempts in running Internet services over broadband. Over time, the systems will become more refined and the underlying business dynamics will make it clearer where development effort would be most effectively applied.

Internet over broadband promises several business advantages—new services and responsive, reliable, and secure service infrastructures—that can lead to new revenue sources. Will it evolve exactly as outlined above? It would be surprising. Will it evolve totally differently? Again, it would be surprising. Whatever the future holds, this discussion has provided a framework for understanding the issues as they are debated and a better intuition for why some approaches succeed more than others.

Application of Telecommunications Technologies in Distance Learning

Ahmed S. Khan
Senior Professor
Department of EET/Telecommunications Management
DeVry Institute of Technology

Abstract

Recent advances in telecommunications technologies (computer networks, satellite communications, fiber-optic systems, the Internet, and so forth) have transformed the modes of learning and teaching. The dissemination of knowledge is no longer confined to constraints of physical premises, with an instructor and a textbook, and it is no longer the only educational resource. The use of telecommunication technologies in distance learning overcomes the barriers of distance and time, and allows students in developed as well as in developing countries to learn in a synchronous or asynchronous manner. This paper presents an overview of various telecommunications technologies that are used in distance education. The discussion includes evolution of distance education, synchronous and asynchronous learning, and characteristics of broadcast television, instructional television fixed service (ITFS), microwave systems, satellite systems, direct broadcast satellite (DBS), cable systems, private fiber, and the Internet.

Introduction

Distance education or learning is a discipline that links people with information through a variety of technologies.[1] Distance learning has been practiced for over 100 years, primarily in the form of correspondence courses. These print-based courses solved the problem of geographic dispersal of students for specialized courses of instruction. In the late 1960s and early 1970s, emerging technologies afforded the development of a new generation of distance learning, aided by open universities, which combined television, radio, and telephone with print.

Today, advances in telecommunications technologies are again changing the face of distance learning. With the explosion of Internet tools over the last few years, distance learning has become more accessible to people whether they live in urban or rural areas. In this time of increasing technological advances, it is no surprise that college enrollment is on the rise. The increasing enrollment, coupled with yearly budget cut requests, opens the door for distance learning.

With the advent of distance-learning technologies, the dissemination of knowledge is no longer confined to the constraints of physical premises. Distance-learning technologies overcome the barrier of distance to allow face-to-face communication between students and teachers from different locations. With rapid technological change in developed countries, educational institutions are challenged with providing increased opportunities without increased budgets. In developing countries, the policy-makers face a dilemma: how to provide educational opportunities to an increased population while considering lack of resources. Distance education offers a solution to the problems of educators in developed as well as developing countries. In developed countries, there are three major motivations for an educational institution to incorporate distance learning in its program offerings:

- to compensate for the lack of specialist courses at its own institutions
- to supplement an existing curriculum to make it more enriched
- to compete with other institutions which offer distance-learning courses via the Internet

For developing countries, distance education offers a unique approach to promote literacy and enhance higher education in a cost-effective way.

Modes of Distance Learning

Distance learning can be employed in a synchronous as well as an asynchronous mode. In a synchronous mode, the student and teacher interact in real time, whereas in an asynchronous mode, the student learns at a convenient time and place.

Advances in computer and communication technologies have greatly contributed to asynchronous access. Asynchronous learning is time- and place-independent; learning takes place at the convenience and pace of the learner. A student can contact a college or teacher via e-mail or engage in discussion with a group through a conferencing system or bulletin board. A learner can participate interactively in a team project with other students that requires problem analysis, discussion, spreadsheet analysis, or report preparation through modern commercial groupware packages. Similarly, lectures can be transmitted through computers, videotape, or CD-ROM. Learning becomes a distributed activity, and participants in these distributed classes access resources and interact asynchronously, more or less at their own convenience.[2]

Asynchronous learning has been categorized into three different levels. They are on-campus, near campus, and far from campus. The on-campus level meets in the traditional classroom, but through computer labs students participate in asynchronous communication through listservs, bulletin boards, and the World Wide Web (WWW). The near campus level (50-60 miles from campus) requires students to meet on campus occasionally for tests, teacher consultation, or labs, but use the Internet or another electronic medium for lectures, class information, and communication. The far from campus level is the "true" asynchronous delivery. This is where the student never sets foot on the campus, and the entire course is delivered through tools such as the Internet, videos, or CD-ROMs. All three levels meet the needs of different learners. One of the benefits of asynchronous learning is that it meets the needs of many "non-traditional" students who would normally not have access to an education due to distance or time constraint.[3]

From Britain to Thailand and Japan to South Africa, distance learning is an important part of national strategies to educate large numbers of people rapidly and efficiently. In the United States, educators are finding distance learning effective not only for outreach to new populations, but also as an important medium for new instructional models. Distance learning, used as a term associated with new technologies offering a full-fledged alternative to classroom education, got its biggest boost internationally with the founding of the British Open University (BOU) in 1969. BOU gained rapid visibility and recognition by broadcasting its video course components weekly throughout the United Kingdom on the BBC network.[4]

Table 1 lists the enrollment for distance learners seeking university education in various developed and developing countries. The addition of distance learners, who receive programming in elementary, secondary, training, and non-credit areas would almost quadruple these numbers.[5]

Many agree that distance learning is the fastest growing instructional pattern in the world. Technologies of delivery, particularly those related to telecommunications, have had a "greening" effect on the distance-learning enterprise. The potential to solve access, cost, time, place, and interactivity considerations that have plagued education since the beginning of time has never been greater.[6]

Telecommunications Technologies

Telecommunications technologies provide opportunities for basic education as well as advanced education for those disadvantaged by time, distance, physical disability, and resources. *Table 2* lists the evolution of telecommunication technologies. *Table 3* lists the breakdown of voice, data, and video technologies used in distance education.[7]

Distance learning blends technological infrastructure and learning experiences. The experiences and applications that can be provided depend on the appropriate infrastructure. Current technologies have favored applications where a central source controls information flow. Emerging networked technologies allow for both centralization and decentralization. While centralized and decentralized approaches each have advantages and disadvantages, networked computer technologies and the distributed-learning applications they enable have not only practical benefits for information management but also attributes that are more congruent with newer educational paradigms based on cognitive approaches to learning.[8]

Broadcast Television

Broadcast TV involves the transmission of video and audio over standard VHF and UHF channels to reach a large number of sites in a limited geographic area (e.g., campus, metropolitan, county, multi-county, and so forth). Most educational programming broadcast over television does not allow real-time interaction with the television instructor but interactively can be designed into live or recorded telecourses using the telephone, computer networks, and Internet tools. This mode of transmission lacks security and confidentiality.

Instructional Television Fixed Service (ITFS)

ITFS refers to a band of microwave frequencies originally set aside by the Federal Communications Commission (FCC) in 1963 exclusively for the transmission of educational and cultural programming. ITFS is similar to broadcast television but is used in a more limited geographic area for simplex point-to-multi-point transmission. ITFS uses omni-directional microwave signals in the 2.5-GHz band to transmit standard 6-MHz video signals to remote locations. An ITFS network can serve as a stand-alone distance education-delivery system, transmitting locally originated programming directly to local schools or cable companies for redistribution through their network. It offers limited or no interaction and a moderate level of security and confidentiality.

Microwave Systems

Terrestrial microwave systems require a line-of-sight transmission between the transmitter and the receiver site. Microwave systems allow simplex or half duplex/duplex point-to-point audio, data, and video transmission. There are two types of point-to-point microwave systems: short haul and long haul. Short-haul systems typically have a range of 5 to 15 miles, suitable for local communications between two schools or campuses. Long-haul systems typically have a range of up to 30 miles between repeaters, depending on transmitter power, terrain, dish size, and receiver sensitivity. Microwave systems offer a high level of interaction, security, and confidentiality.

Satellite Communications

Satellite communication involves the transmission of a broadcast signal from an earth station via an uplink signal to a geosynchronous satellite, where the signal is processed by the transponder and sent, via a downlink signal, to a receiving dish antenna. Satellite communication is used to cover a large number of sites over a wide geographic area (e.g., statewide,

TABLE 1

Major Open Universities in the Developed and Developing Countries

Institution	County	Year established	Enrollment
University of South Africa	South Africa	1951	50,000
Open University	United Kingdom	1969	50,000
Universidad Nacional	Spain	1972	83,000
Fernuniversitat	Germany	1974	37,000
Open University of Israel	Israel	1974	12,000
Allama Iqbal Open University	Pakistan	1974	150,000
Athabasca University	Canada	1975	10,000
University Nacional Abierta	Venezuela	1977	29,000
University Est A Distancia	Costa Rica	1977	11,000
Sukhothai Thammathirat OU	Thailand	1978	200,000
Central Radio and TV University	China	1978	1,000,000
Open University of Sri Lanka	Sri Lanka	1981	18,000
Open Universiteit	Netherlands	1981	33,000
Andrha Pradesh Open University	India	1981	41,000
Korean Air and Correspondence University	South Korea	1982	300,000
University of the Air of Japan	Japan	1983	22,000
Universitas Terbuka	Indonesia	1984	70,000
Indira Ghandi Open University	India	1986	30,000
National Open University of Taiwan	Taiwan	1986	48,000
Al-Quds Open University	Jordan	1986	not available
Universid ade Aberta	Portugal	1988	3,800
Open University of Bangladesh	Bangladesh	1992	not available
Open University of Poland	Poland	Proposed	
Open University of France	France	Proposed	

TABLE 2

Evolution of Telecommunication Technologies

Year	Technological Development?
1832	Telegraph
1875	Telephone
1895	Radio
1945	Audiotape
1953	Broadcast Television
1960	Videotape
1960	Audio Teleconferencing
1965	Cable Television
1975	Computer Assisted Instruction
1980	Audiographic teleconferencing
1980	Satellite Delivery
1980	Facsimile
1980	Videoconferencing
1984	Videodisc
1985	CD-ROM Compact Disc
1988	Compressed Video
1989	Multimedia
1990s	Lightwave systems, LANs, HDTV, Internet

TABLE 3

Telecommunication Technologies in Distance Learning

Technology	

Voice	Telephone
	Radio
	    Short-wave
	    AM
	    FM
	Audiotapes

Data	LANs
	WANs
	Internet
	    Computer-assisted instructions (CAI)
	    Computer-mediated education (CME)
	    Computer-managed instruction (CMI)

Video	Pre-produced videotapes Videoconferencing Broadcast television ITFS Satellite System DBS Internet

nation-wide, and continent-wide) for simplex point-to-multipoint audio and video transmission. Analog satellite systems offer low levels of security and confidentiality, and digital systems offer high levels of security and confidentiality. Satellite systems offer limited or no interaction.

Direct Broadcast Satellite (DBS)

These high-powered satellites transmit programming directly to the general public. The received dish antennas used in DBS systems are very small (<1 meter). DBS systems allow programmers to beam educational programming directly to homebound students, providing an alternative to over-the-air broadcast or cable television. DBS technology employs data-compression technologies that enhance the efficiency of video channels.

Televised education offers the following advantages:[9]

- Televised courses can actually enhance instruction by allowing instructors to present material in novel ways and thus to increase interest in it.
- Televised courses provide access to education for those who might not otherwise have it.
- Televised courses can help the instructor become a better teacher.
- Televised courses are often better organized and more highly developed than traditional courses.
- Televised courses offer an opportunity for faculty members to learn from their students as well as teach them.

Satellite communications is one of the most cost-effective modes to promote literacy and enhance higher education in developing countries. *Table 4* lists the typical cost for leasing a video channel for satellites that provides large footprints in Asia and Africa.

Cable Systems

Cable-television systems use coaxial and fiber-optic cable to distribute video channels to local subscribers. Programming is received from local broadcast channels and national programming services at the cable "head end" and is sent out over the cable in a tree configuration. A cable head end can receive many types of signals, such as satellite or microwave transmission, which can then be retransmitted to schools over the normal cable system.

Cable-television systems are primarily one-way (simplex) broadcast (point-to-multipoint) type transmission systems. Many systems also have a limited number of reverse channels, providing some measure of two-way interactivity. There is a high-level inter-system security and confidentiality.

Private Fiber

Optical fiber is used for full-duplex point-to-point audio and video transmission. It allows a high degree of interaction and offers a high level of security and confidentiality. It also has the highest information-carrying capacity. The bandwidth for mono-mode step index fiber is 10-100 Gbps; multimode step index fiber is around 200 Mbps.

TABLE 4

Typical Cost for Leasing a Video Channel for Satellites Providing Large Footprints in Asia and Africa

Satellite	Footprint coverage area	Typical hourly cost for leasing a video channel
Panamsat	South Asia	$1850
TDRS-5	South Asia	$1680
Chinasat 1	Southeast Asia	$990
Asiasat 2	Southeast Asia	$990
Measat	Southeast Asia	$990
Palapa-B2	South Asia	$885
Intelsat-k	Western Africa	$2000
GESTAR-4	Western Africa	$2000
Inmarsat-3	Western Africa	$2000
Intelsat	South Africa	$1000
Orion 2	South Africa	$1000
Panamsat PAS3	South Africa	$1000

Public Telephone Service

Public telephone service is used over a limited to wide geographic area (local, regional, national, and international) for full-duplex point-to-point or point-to-multipoint audio transmission. It is used in conjunction with other technologies (television, satellite, etc.) to provide a feedback channel for students at remote sites to interact with the instructor and other sites. *Table 5* compares the characteristics of distance-learning technologies.

The Internet

The Internet offers full-duplex point-to-point and point-to-multipoint audio, video, and data transmission over a wide geographic area (regional, national, and international). It allows a high degree of interaction. Generally, it has a low level of security and confidentiality but, with the use of data encryption, it has a high level of security and confidentiality.

With a computer, minimal software, and an Internet connection, people anywhere in the world can access on-line sites and programs. In addition, these programs can be accessed almost instantaneously through a user-friendly "point and click" interface, using almost any type of computer without producing and distributing printed materials, floppy disks, or CD-ROMs and with no training delivery cost. Furthermore, training can be delivered to a potentially unlimited audience, updates can be easily made on-line, and programs can link to a vast collection of other on-line resources.[10]

Although the Internet can be accessed worldwide, current access for many people is still limited. Even when access is available, bandwidth for accessing the Internet is frequently low. Low bandwidth severely limits capabilities such as on-line interactive multimedia training. Other disadvantages are that students need a moderate degree of computer literacy; sensitive/classified training requires additional security measures; and the high level of current Internet "hype" makes it difficult to determine what is practical for the near future.[11]

In addition to these technologies, a number of high-speed non-switched and switched lines could be incorporated in wide-area networking strategies in distance learning. *Table 6* lists various switched and non-switched services provided by U.S. carriers.

A recent report entitled "Distance Education in Higher Education Institutions" issued by the National Center for Education Statistics highlights the following facts:[12]

- By fall 1998, 90 percent of all institutions with 10,000 or more students and 85 percent of those with enrollments of 3,000 to 10,000 expect to offer at least some distance education courses. Among the smallest institutions (those with fewer than 3,000 students), only 44 percent plan to offer distance education courses.

- More than 750,000 students were enrolled in distance education courses in 1994-95. That year, some 3,430 students received degrees exclusively through distance education.

- More students enroll in distance education courses through public two-year institutions than through any other type of institutions.

In terms of the type of technology used to deliver courses, 57 percent of the institutions offering distance education courses in 1995 used two-way interactive video, and 52 percent used one-way prerecorded video (institutions frequently used more than one type of distance-learning technology). Twenty-four percent used two-way audio and one-way video, 14 percent used two-way on-line interactions, 11 percent used two-way audio, 10 percent used one-way audio, and 9 percent used one-way live way.

A key point to note, however, is that because of the many logistical factors involved, distance education courses typically require considerable investments of money and time. Distance education courses can be cost-effective when a large number

TABLE 5

Comparison of Distance-Learning Technologies

Technology	Characteristics
Broadcast Television	Simplex mode of transmission One site to multi-site video and audio transmission (point-to-multipoint transmission) Limited or no interaction Public system Lack of security or confidentiality
Instructional Television Fixed Site (ITFS)	Similar to broadcast TV Semi-public system Moderate level of security and confidentiality
Broadband Cable	Similar to broadcast TV Public and private systems High degree of interaction possible High level of inter-system security and confidentiality
Microwave	Point-to-point audio and video transmission over limited area (one site to another) FDX mode of transmission High level of interaction Private system High level of security and confidentiality
Satellite	Simplex mode of transmission Point-to-multipoint audio and video transmission over wide area Limited or no interaction Private or public system Low level of security and confidentiality for older analog systems; higher level of security and confidentiality for digital systems
Private Fiber	Point-to-multipoint transmission over limited area FDX audio/video transmission High degree of interaction possible Private system High level of security and confidentiality
Public Switched Digital Service	Point-to-point and point-to-multipoint audio and video transmission over short and long distances Moderate to high degree of interaction possible Public system High level of security and confidentiality
Public Telephone Service	Point-to-point and point-to-multipoint audio transmission over short and long distances. Used in conjunction with satellite and other simplex technologies to provide a way for remote site to interact with instructor and other sites
Public Packet Network Services	Point-to-point SPX transmission Used in conjunction with other technologies to provide a way for students to interact with the instructor

Source: http\\www.oit.itd.umich.edu/reports/DistanceLearn/sect.html

TABLE 6

High-Speed Wide-Area Networks

Non-switched (leased)	
Analog	4.8-19.2 kbps
Digital Data Service (DDS)	2.4-56 kbps
T-1	1.54 Mbps
T-3	44.736 Mbps
Frame Relay	1.54 Mbps - 44.736 Mbps
Synchronous Optical Network (SONET)	51.84 Mbps - 2.488 Gbps
Switched	
Dial-up / modem	1.2 - 28 kbps
X.25 packet switching	2.4 - 56 kbps
Integrated Services Digital Network (ISDN)	64 kbps - 1.544 Mbps
Frame Relay	1.54 Mbps - 44.736 Mbps
Switched multimegabit data service (SMDS)	1.54 Mbps, 44.736 Mbps
Asynchronous Transfer Mode (ATM)	25 Mbps - 155 Mbps
B-ISDN (broadband ISDN)	155 Mbps, 600 Mbps

of students take the courses over the long run. It is important to realize that such courses are costly to produce and deliver. The costs rise even more when the courses are presented with advanced one-way or two-way real-time video.[13]

The incorporation of new telecommunications technologies in distance education has abolished the barriers of location and time. However, the success of distance education in higher education depends on:[14]

- the availability of high-quality, full-motion video via the Internet
- the electronic accessibility at a reasonable cost of content currently available only in print format
- development and documentation of effective teaching / learning processes that rely on advanced technology
- the ability of distance learning to comply with an accreditation process for engineering education that is outcome-based

Conclusion

Distance education has come far since its humble beginnings as correspondence courses conducted by mail. To excel in the 21st century, higher education must undergo a paradigm shift. It must be transformed from an environment and culture that defines learning as a classroom process shaped by brick-and-mortar facilities and faculty-centered activities to an environment defined by "learner-centered" processes and shaped by telecommunications networks with universal access to subject content material, learner support services, and technology-literate resource personnel.[15] The use of telecommunication technologies in distance learning will promote literacy and enhance higher education. Thus, distance learning will narrow the techno-economic gap between the developed and developing world of the 21st century.

References

1 *www.fsu.edu/~lis/distance/brochure.htm#program*

2 Mayadas, F. and P. Alfred (1996) *Alred P. Sloan Foundation Home Page.* (Online). Available at *www.sloan.org/Education/ALN.new.htm.*

3 *Linking for Learning: A New Course for Education,* Office of Technology Assessment (OTA), U.S. Congress, Washington, D.C, 1989.

4 Granger, D., "Open Universities: Closing the Distance to Learning, Change," *The Magazine of Higher Learning,* 22 (4), pp. 42-50.

5 Brown, B., and Y. Brown, "Distance Education Around the World," *Distance Education Strategies and Tools,* edited by Barry Wills, Educational Technology Publications, NJ, 1994.

6 IBID.

7 Willis, B., *Distance Education Strategies and Tools,* Educational Technology Publications, Englewood Cliffs, NJ, 1994.

8 Locatis, C. and M. Weisberg, "Distributed Learning and the Internet," *Contemporary Education,* vol. 68, no. 2, Winter 1997.

9 Richard, Larry G., "Lights, Camera, Teach," *ASEE Prism,* February 1997, p. 26.

10 Collis, B., *Tele-Learning in a Digital World: The Future of the Distance Learning,* International Thompson, London, 1996.

11 IBID.

12 *ASEE Prism,* February 1998.

13 Martin, B., P. Moskal, N. Foshee, and L. Morse, "So You Want to Develop a Distance Education Course?" *ASEE Prism,* February 1997, p. 18.

14 *ASEE Prism,* February 1997, p. 4.

15 Dubols, Jacques R., "Going the Distance, A National Distance Learning Initiative," *Adult Learning,* September / October 1996.

Connecting Carrier and Customer: Network Management Over the Internet

Steve Kim
Vice President—Engineering
Clear Communications Corporation

The Internet's popularity has come about not only because of easy access, but also because of the potential for instant gratification on-line. Businesses can get competitive data immediately, and consumers can look at products and order them without entering a store or talking to a customer-service representative on the telephone.

This potential for instant gratification is one of the most important things that a company needs to provide in order to retain customers and gain customer loyalty. For example, the proliferation of gasoline-pump credit-card systems produced brand loyalty for the retailers who installed them first.

Customer Service Concerns

Part of the customer-care challenge for service providers arises from the fact that they are such large organizations with multiple points of contact. Although it is usually possible to work with a service provider via telephone, at times there is quite a bit of telephone tag involved. All the information needed by the customer is not immediately available, and varying degrees of information will be provided, depending on the given service representative a person speaks with. Some of the information is useful, some is not, and some is not consistent with what the customer has heard before. The result of this is reduced confidence in the service provider, slower introduction of new services, and higher customer churn. When customers get a bad feeling, they will go to one of the company's competitors.

For example, when your air conditioning or heating breaks down in your home, you will probably go through the phone book to find a repair service. But if the serviceman says he cannot come out until late in the day or perhaps three days from now, you will probably look for another company that can send someone sooner. This illustrates how quick response and instant gratification in regards to a customer-service problem is very important to customer loyalty: you will continue to use that heating-repair service as long as they continue to do a good job in providing service.

Banking is one industry which has used customer network management as a competitive advantage. Automatic teller machines provide access to cash, allow customers to deposit money, and provide other services quickly and easily. American Airline's Saber reservation system and Federal Express's parcel tracking service are examples of systems that allow for on-line access to information very quickly without having to talk to a company representative. These companies are in the forefront of using the Internet because it provides a step forward in instant access and gratification.

Who will win the customer loyalty game? Giving customers information in a very timely and easy fashion via customer network management will play a tremendous role in enabling service providers to retain customers and build customer loyalty. Service providers who can provide this instant gratification will have a tremendous advantage in winning customers and market share. There are a number of players in this game now—interexchange carriers, incumbent local exchange carriers, competitive local exchange carriers, cable operators, competitive access providers, and others. The main challenge for companies such as these is to partner with their customers to provide business solutions.

Internet Solutions

Through customer network management (CNM) and the Internet, service providers can gain a competitive advantage by presenting customers with a more consistent and appealing interface. For example, an order-entry application could enable customers to order new products and services without having to talk to a service representative or an operator. Other types of information that could be relayed in this way include network status, help-desk assistance, trouble reports, and service configuration. Any form of service report could be produced on-line, which would enable a customer to avoid waiting in a queue to talk to a service provider.

One of the main challenges to customer network management, particularly over the Internet, is partitioning customer information to ensure confidentiality. Customers need access

FIGURE 1

Management of CNM Information Flow

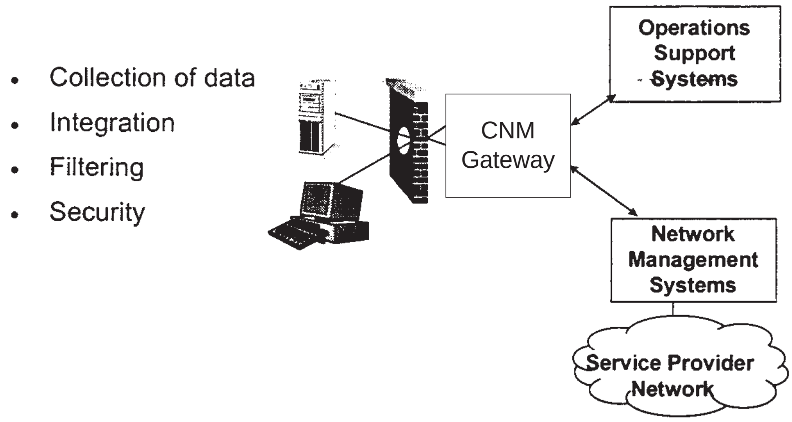

to their own information and the assurance that nobody else can access that information. *Figure 1* shows one possible CNM gateway solution. Information must be partitioned by the customer, and management of the CNM information flow must provide for collection of data, integration, filtering, and security.

There are several ways to provide customer network management capabilities (see *Figure 2*). The Internet is one method, and the others are via remote workstations or via machine-to-machine interfaces between the customer's management system and the carrier's CNM gateway. Nevertheless, the option that will take off in the future will be the Internet because of the factors discussed already: easy access and instant gratification.

A customer network management system must be able to meet the special needs of wholesale and retail customers. For wholesale customers, the system must allow the customer to control the unbundled network with its own network management system (NMS) (see *Figure 3*). Customers must be able to manage and understand their virtual access network, which involves inventory management, provisioning, fault management, and performance management. This is all part of providing a partitioned environment so that a wholesale customer can perform these functions in a secure, partitioned manner. In contrast, retail customers are more interested in the user interface that provides them access to the services that they are purchasing from wholesalers (see *Figure 4*). This interface must be consistent, and the Internet provides that advantage. It provides a graphical view of the network and an industry-standard interface (via browsers) that is accessible from any location.

Operational Features

Part of an effective customer network management system is answering some of the day-to-day questions that customers

FIGURE 2

CNM Provider Capabilites

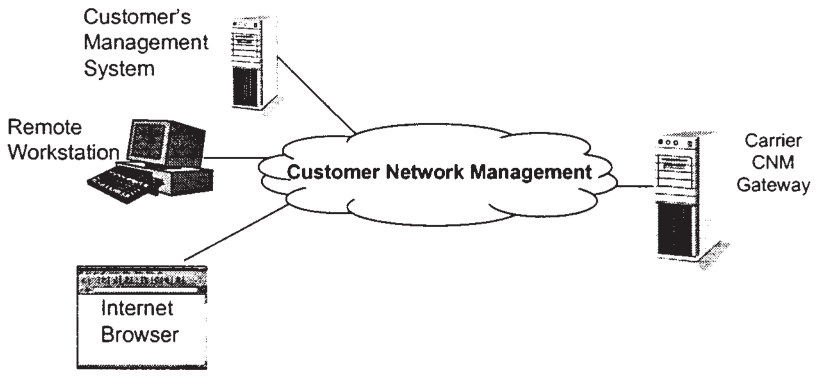

FIGURE 3

Management of Virtual Access Network

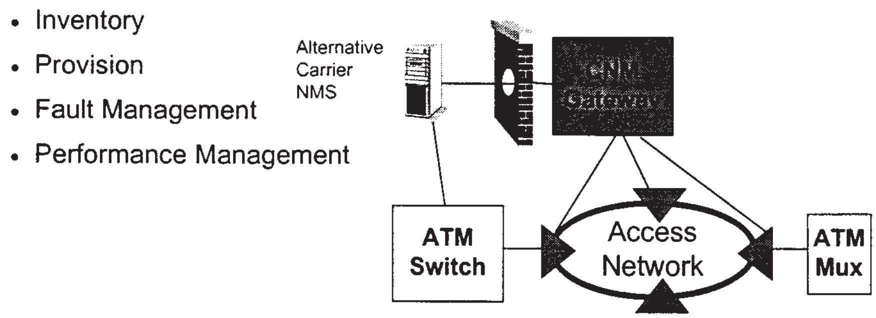

will have, in a quick and easy manner. Some of these questions include:

- What is my network?
- What are the components and elements that make up my network?
- What circuits make up my network and how does it operate?
- How is my network performing today? Is it in line with the service agreements that were signed with my service provider?
- How do I get help when I need it? Do I have to call a help desk, or can I enter a trouble ticket on-line?
- How can I see where my trouble ticket is at any given time? How much longer is my problem going to last?
- Can I order a product on-line? If so, how long will it take to arrive?

Fault management is an important area, and it too can be off-loaded into a customer network management application. Fault management involves management per circuit, including service ID, description, date / time, as well as the affected customer and network. Providing access to alarms and events can be enhanced by root cause analysis (RCA) so that a customer can identify the actual alarm that is causing a problem, rather than looking through all the alarms that are in the network (see *Figure 5*). Status information that the system should provide includes an inventory of services and circuits that are part of a particular customer's network; it should also provide the ability to see how the network is performing by quickly identifying alerts and status indicators.

Sometimes network operators can become inundated with data, but it is often not the information they are looking for. Internet broadcasting can address this problem by providing a mechanism which customizes the information that is sent according to what the customer actually wants to see. Push technology, which is currently offered by companies such as PointCast and Marimba, can be used to automatically deliver information to the customer about their service orders, trouble tickets, and network performance status.

FIGURE 4

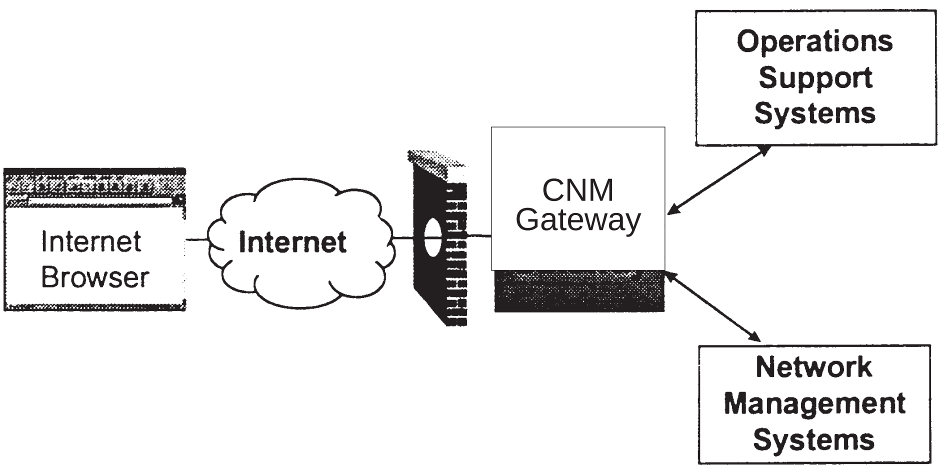

FIGURE 5
Fault Management

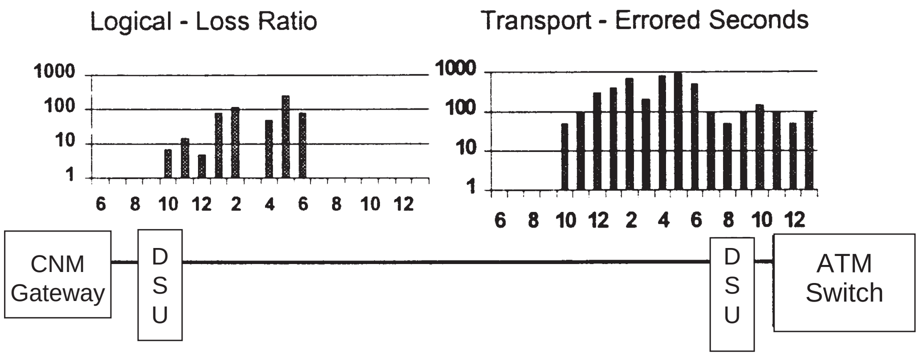

Technical support is another important application of CNM. On-line access to a help-desk technician, answers to technical questions, open trouble reports, and tracking of problem resolution are very important applications for CNM. Information about the historical performance of the network is also valuable to have available on-line via a CNM system. This includes consolidated network-performance data, individual end-to-end service, performance-trending information (daily, monthly, and yearly), and problem-circuit indications. Providing these reports on-line enables customers to quickly determine whether their networks are in compliance with their service level agreement (SLA). The system should also allow customers to measure their network quality of service, which involves the flexible reporting of network and circuit reliability, availability, utilization, and performance.

Integrating the service order system and service activation with the carrier's CNM gateway enables customers to get a new service when they need it. This includes the following functions:

- request service activation
- confirm service configuration
- track status as circuit is installed
- final acceptance of service
- enable the service when it is ready to use
- suspend or disconnect the active service

Underlying all these capabilities, however, is the requirement for network security. An effective security access system will include name and password authentication, as well as adequate Internet security standards (i.e., a secure socket-layer solution). In addition, encryption of transported data and optional TCP/IP address specification is needed.

Summary

Customer network management needs to be thought of as a strategic value-added service that is a very important component for securing customer loyalty, market share, and customer retention. Wholesale customers need electronic interfaces to network elements and support systems, and retail customers need the graphical interface and Internet support in accessing information. In summary, customer network management is genuinely moving from a "nice-to-have" service to a "must have" capability for service providers in order for them to compete.

The Internet: The Start of the Same Old Billing Problems

David Lamb
Product Manager
Axiom Inc.

Abstract

This paper examines the development of the Internet from a historical, technical, business, and customer perspective. An evaluation of funding models is provided, as well as an explanation of why usage metering will be a part of the Internet funding model. Finally, an analysis of the issues and potential solutions to usage-based event collection is outlined with a conclusion that the telephony industry has solved many of the associated problems and, in some cases, has the appropriate infrastructure (mediation system) to resolve event-based billing for the Internet.

The Internet Evolution

In 1969, a U.S. Defense Department agency, Advanced Research Project Agency (ARPA), created the foundation of the Internet as a military application. Known as ARPANET, it was the first wide-area, packet-switching network that allowed sharing of information. In 1983, the U.S. government separated the military and civilian (mostly university-related) applications of ARPANET.

In 1986, the National Science Foundation (NSF) created NSFnet—a new rival, non-military network to share information—that subsequently replaced ARPANET in 1990. The NSFnet created a wide-area network (WAN) between five large data centers and a significant number of universities. The management of the network was a cooperative agreement between Merit Network Inc., MCI, and IBM, which eventually was formalized as a joint venture—Advanced Network and Services (ANS). MCI was responsible for the physical links, and IBM provided support for the routers.

In 1995, the NSFnet replaced backbone services provided competitively by a small group of commercial providers including MCI, PSINet, UUNET (now WorldCom), ANS / AOL, and Sprint to name the major players. These providers interconnected at a series of network access points (NAPs) to create the commercial Internet of today.

Before 1990, when NSFnet replaced ARPANET, the network-based services were limited to an elite research-based government / university population. Starting in 1990, however, the Internet became a commercial vehicle that could be accessed by more and more people. The Internet as a technology has been available since the late 1960s, but as a commercial service has only been provided since 1990. The assumption that the Internet became available in 1990 has a direct linkage as to when a predominant funding model will be defined, as explained later in the paper.

The Internet Operational Architecture

The Internet's technical architecture can be segmented in a number of ways:

- Internet access
- Internet service provider
- backbone network

Internet access is defined as a point of access to an Internet service provider (ISP), which is provided by a telecommunications company via modem, xDSL, 56-k, or T1 (or above) leased line. For the most part, residential customers will use modem and xDSL access, while business customers will use fixed network connectivity. In addition, cable companies are now offering cable modem access for residential customers via cable servers with the potential for far faster Internet access than the existing POTS network can support.

An ISP will house a server (anything from a Pentium PC upward) and provide the routing and access to the rest of the networked computers that constitute the Internet. Typically, the ISP will provide "home page" services and other content services on the server. The ISP will be connected to the backbone network via 56-k leased line or above network to a NAP.

The backbone network consists primarily of packet and ATM switches with networked computers routing TCP / IP packets of data to the appropriate "IP" address. This network is consistently being upgraded to handle the growing number of customers and increasing data-rich environment.

The Internet Business Model

A fundamental question for the Internet business model is a definition of what the Internet is. Specifically, the Internet needs to be defined as either a content service, which means a television or radio business model based on advertising revenues, or as a telecommunication service, which incorporates usage pricing. The Internet's amazing growth is equal to the growth of radio, television, and telephones, all of which changed society considerably. If we look at how long it took to establish the enduring business model for the "old" media, one should expect the Internet to take about 20 years to establish its funding model. However, based on historical review, the Internet "start date" is arbitrary at best, but as outlined, 1990 is the most appropriate year.

The ISP business model is and will be defined by the adult and gambling services as services that are pioneers by the very nature of their business. A March 1997 article in *Inter@ctive Week* sized the adult element of the Internet at about $1 billion in revenue per year with approximately 10,000 sites. With such an attractive revenue stream and high level of competition, innovation is also being applied to the funding model, as well as the adaptation of technology to provide service. Some large sites are taking millions of hits daily and generating more than $1 million per month in revenue for the service operator. Technical innovation from this sector has included push technologies, video streaming, and sound over the Internet. The original business wisdom was for a flat-rate charge per month, but that changed in 1996 when home pages started financing banner ads. With 250,000 visits per day, banner ad rates rose from $1,500 per month to $4,500 in only three months. Some home pages are now evaluating pricing per click, a primitive usage-based model. ISPs are using a combination of funding models to generate their revenue, and it is unclear whether the entertainment or the telecommunication model will be adopted.

However, in the long term, end users of Internet services will not be interested in multiple sources for billing. Just as MCI has packaged telecommunication services under the MCIOne brand, the large players in the Internet market (including the telecommunication companies) will offer an integrated / bundled basket of products with an associated pricing model that will include Internet access, ISP services, and network utilization. These combined services will be based on assembly of a number of usage-based events. Pricing may be below cost, due to revenues from advertising, discounting the end customer service price. However, usage aggregation and analysis will be essential for the long-term success of ISPs as they try to balance investment in infrastructure with resource utilization and revenue-generating capabilities.

Internet Usage-Based Pricing

Internet usage-based pricing will be adopted, if only between the Internet layers of service / access. Inter-administration accounting will be part of the funding model for the network / ISP wholesalers. The end customer may not see any of the usage events as the service provider may have bundled all the associated costs and offset them against advertising revenues. However, the service provider will be invoiced by all the organizations responsible for the delivery of service to the ISP's customer—very much like the existing telecommunications model. The end customer wants to have one integrated bill that is easy to understand. Customers cannot be expected to equate a usage rate per data packet with the service they are utilizing on the Internet.

Now that the case for usage event processing has been established, consider the usage events that will be created by an Internet business event:

- xDSL access will create data packet events.
- POTS utilization will create a connection event.
- ISP access will create CPU events.
- Backbone network access will create data packet events.
- Server access will create another CPU event.
- Entertainment service on server will create a usage event.
- Access back to the originator will create data packet events.

The number of events that require collection and pricing for a simple Internet service is considerable. Unlike many telephony services, all the billable data will not be in the first segment of the event, and for bundled service providers, a form of aggregation will be required to produce a billable event. The billable event will consist of the following vectors: duration and time of access, network resources utilized (data packets), CPU utilized (processing cycles), and content. Telecommunication companies will be providing many of the service layers within the Internet and will, for their end customers, need to aggregate the usage events from many diverse application and network elements.

Furthermore, this requirement will extend to real-time processing, possibly to provide a rate prior to the initiation of the event. Put simply, customers will shop around the many similar service offerings on the Net and try to find the best price for the specific event they want to initiate. Thus, the standard pricing question for any commercial transaction will apply: How much is it? The service provider will need to quickly aggregate the known price components for that specific service and provide an answer within seconds. Finally, to add further complexity, services must have the ability to include billing for other parties using the same service. For example, when a user decides to play his or her friend at chess via the Internet, and it is agreed that the loser will accept the charge for the chess service on his / her ISP account.

Barriers to Usage-Based Pricing

Among many of the long-term Internet users, there is a near-religious belief that the Internet is a public service that should remain free. Today's average user has adopted this attitude. It will be difficult to change this perception. However, to appreciate the potential for non-usage-sensitive pricing, one has only to review some of the capacity problems that AOL experienced where usage was not an issue, due to flat-rate pricing, and thus customers were not logging off. As with all markets, there will be tiering of service levels in terms of guaranteed access, response times, and problem resolution. However, the price associated with these upper level tiers will be commensurate with the quality of service.

The definition of a business event as mentioned earlier in the paper requires considerable analysis and socialization. Charging by the packet has no relevance to the end customer as it is impossible to assess final price of a service with price per packet. Thus, pricing must relate to an event the customer can identify with. As mentioned, the requirement is to develop the appropriate tools to assemble events to a meaningful customer event. This is not simple when there is such a broad possibility of components that the customer can assemble to define the service that is required.

The investment required to capture, assemble, and rate events will be considerable. The business case for the funding of the necessary infrastructure has not been established. Flat-rate pricing or advertising revenues are simpler to administer and require significantly less investment to realize.

Finally, the rapid technology change that is enabling all imaginable (and many unimaginable) services to be provided is glued together via a set of loose standards. To this effect, there is no standard for presentation or format. Thus, the requirement for assembly of transaction events will also require data normalization, an additional processing overhead, and another variable for the ISP to manage.

Internet Mediation Requirements

Within the world of telephony, the evolution and necessity of standards for service delivery ensured seamless delivery of service to the end customer. For the most part, the same is true on the Internet. However, interoperability between network components and service elements has and will be a tremendous issue both for telephony and the Internet. Telephony is resolving this via the International Telecommunications Union (ITU) and Telecommunications Management Network (TMN) standards. In the meantime, most large telephone companies are using mediation systems to collect data from the many network elements in their network via many different protocols and then normalizing the data and protocol for downstream rating and fraud applications. Like the Internet, telephony is developing products based around call processing, including advanced intelligent network products, which often require event assembly due to their complexity. These requirements are moving into the real-time world as the complexity of telephony products increases and the telecommunication market becomes more competitive.

Thus, the requirements for Internet-usage event mediation will be the capability to assemble usage (or possible usage events) into billable events that the customer can understand and to do this on a "what-if " basis, due to the multicomponent and assembly capability of the Internet. Given the need to perform what-if pricing, the mediation system must be able to perform the assembly in near real time. As an additional twist on pricing presentation above normal telephony billing, some services will require that the pricing runs as the service runs. An example of this would be a game where in the bottom left-hand corner a price display shows how much money you are spending as the game runs. The Net, which is display-oriented, can provide this representation of invoice.

The lack of standards on the Internet will require collection via whatever protocol is available as well as data normalization, since usage data will be in any and every format. Once the data has been collected, it will require distribution to a number of downstream applications including an inter-administration accounting system, a rating and invoicing system, a fraud-management system, traffic management (utilization) system, and, if a wholesaler, back to the retailer of services.

The rate of change on the Internet is driven by technical possibility and the fact that there is no restraining governing body. This will require that the mediation system is inherently flexible and easy to change / use. The functional requirements defined for this mediation system will drive the processing requirements. It is unlikely that one platform will be able to provide the volume processing and real-time responses required to meet this requirement. Thus, much of the processing (such as the services it supports) will be distributed over the network. Due to the nature of the service offerings, the mediation system will also need to activate services for the customer if the customer wishes to play a game against a non-customer via the Internet, or if the specific service required is restricted in some way. Finally, to manage this distributed architecture, a network manager with the appropriate ease of use (and probably a Web browser) will be required to provide a management center for the solution. *Figure 1* represents a typical implementation.

Internet Mediation Architecture

As mentioned earlier, telephone companies have been resolving interoperability problems in billing for many years. As a statement of requirements, Bellcore's AMADNS GR 1343 CORE document details a mediation infrastructure that will meet both telephony and Internet requirements. Thus, many telephone companies that have invested in flexible LAN-based distributed mediation systems will be well-positioned to utilize the same infrastructure for event capture and analysis from the Internet, reducing the requirement for additional investment.

Conclusions

It will be a number of years before a predominant funding model is established on the Internet. In all probability, it will be a mixture of flat-rate pricing, usage-based pricing, and advertising revenues. However, the Internet, like telephony, is a conglomerate of network infrastructures and service providers that will need to account for resource utilization between the different entities and the associated financial settlements.

The mediation requirements for the medium term will increase (not decrease) as component and service complexity increases with few governing standards. Thus, data collection and normalization in near real time will increase in proportion with the complexity and variety of new services. Like telephony standards, governing formats / interfaces will eventually emerge similar to TMN.

The end customer will require a business event that makes sense in terms of service utilization, value, and financial exposure. Pricing solely based on data packets or CPU cycles will have little relevance to the customer playing interactive war games. Furthermore, the mix-and-match nature of the Internet will require that pricing be able to perform

FIGURE 1

AMADNS Mediation Architecture

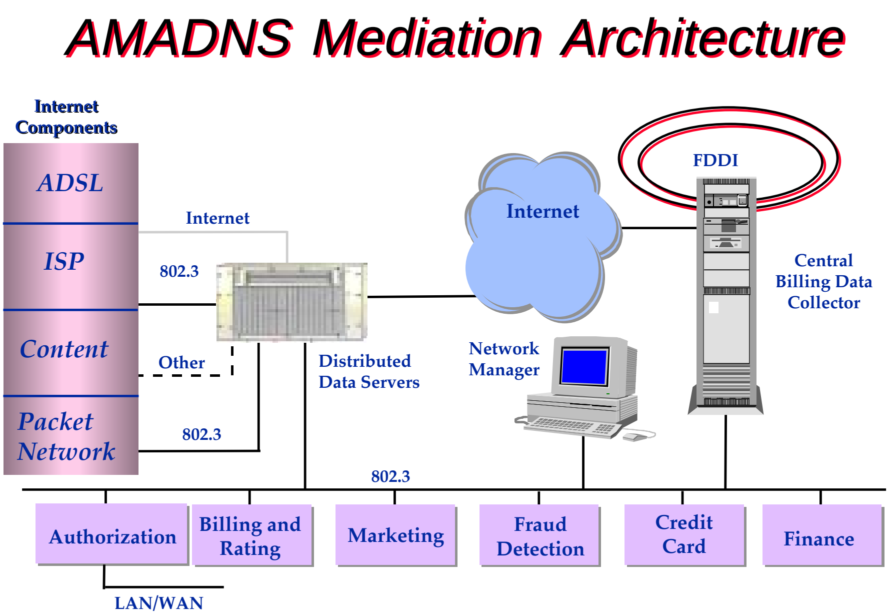

what-if analyses against specific customer assembly of service components.

Many telecommunication companies are already deeply involved in many of the layers of the Internet from core network to ISP to content services. In the future, when a usage-based pricing model is adopted, these services will be bundled into one service offering for the customer. As many of these services are an extension of existing telephony, much of the infrastructure exists to support the bundled Internet business. A key component will be the event mediation systems as long as they have been designed to meet, at a minimum, the requirements outlined in the Bellcore AMADNS GR 1343 CORE.

Internet Impact on Current Network Infrastructure

Lynnwood Lee
Manager, Network Design
Sprint

The Internet has had a substantial impact on current network infrastructures. Local exchange carriers are facing a situation in which their networks have been carrying traffic for a long time, and all of a sudden they have much more traffic than they used to have.

This paper will examine six key issues in assessing the impact that the Internet has had on plain old telephone service (POTS). One issue is how Internet service providers are affecting other customers' quality of service. The next are the issues faced by Internet service providers (ISPs) themselves in trying to accommodate ever-increasing numbers of customers. The switch/remote impact of ISPs will be considered—that is, what happens away from the central office and closer to actual customers—as well as the impact of the Internet on extended area service (EAS) networks and how increased Internet traffic affects trunking for all telephone services. The fifth topic to be addressed is the actual impact on traffic, illustrated by a study of traffic data. Finally, the issue of how to provision for Internet use will be discussed.

POTS and Internet Customer Issues

Before digital-to-central-office technology became standard in the industry, slow dial tones were one of the key customer complaints. After being almost eliminated, the problem has returned with the onslaught of Internet users. With many more users consuming more time on-line, ISPs are causing major problems, including terminating-switch blockage, EAS trunk-network blockage, and busy lines connecting to calling ISPs.

To provide some perspective, there are approximately 690 million phone lines working worldwide. Of those, 172 million are in the United States. Presently, approximately 30 percent of the homes in the United States have personal computers, and of that less than 10 percent have Internet access. By the year 2000, roughly 34 percent of homes nationwide shall have Internet access. Somehow, in the next two to three years, local exchange carriers will have to figure out how to switch all that traffic.

Internet Service Provider Issues

On the other end of the line, ISPs are facing many of the same issues as local telephone companies: They want to be able to provide service to their customers. Fortunately or unfortunately, their service depends upon the public switched telephone network (PSTN) to connect their customers to their ISP servers and routers. This can be a win-win situation if LECs and ISPs work together to size the ISP's lines properly, regulate their connectivity to the public switched network, and enable their customers to get on-line. This is a big issue, one that will cause problems in the future.

If an LEC can work with ISPs, it can come up with better solutions than just line-ended modems. Obtaining a faster connection will be a key issue in the future, but it is also important now: modems are rated at 28.8 kbps and higher, but local phone service is tariffed at 9.6 kbps. Local companies will have to condition their outside plants in order to carry faster speeds and come up to par.

ISP Growth

A number of ISPs started up in residential garages or basements with 10 or 15 lines. They grew to 50–100 lines and moved to commercial offices. For a local exchange carrier, this presents a number of headaches. If the typical residential customer line is two cable pair, how does the LEC get 15 or 20 lines into a garage ISP? What about a strip mall?

Given this situation, the LEC is forced to build an outside plant. Hopefully, the ISP can offer the LEC lines of back-to-back channel banks over a T1 connection. That is not too difficult, but the same issues crop up when the ISP outgrows its space. Sprint is encouraging ISPs to move into some of its own central offices, eliminating a lot of outside plant and enabling them to connect directly into the switch. This can provide tremendous benefits to both parties.

Switch/Remote Impact

Figure 1 shows the Nortel DMS switch architecture. Right now, switch resources can be shared at a 6:1 concentration ratio for the average POTS customer, but that ratio no longer applies when high-speed and Internet connections come into play (the average CCS per Internet customer is 18 CCS/line, whereas the average CCS per POTS customer is 2.4 CCS/line). In a system where there might be 640 lines out in a line-concentrating module, high Internet usage can reduce

FIGURE 1

LCM Choke Points

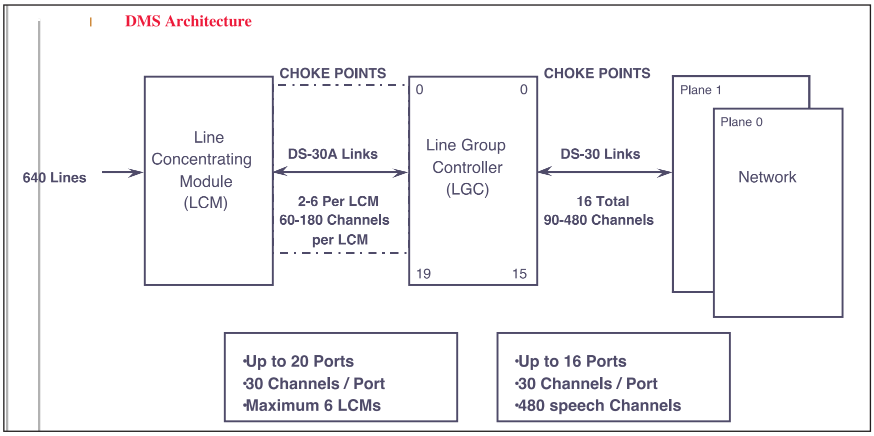

the module to only handling 180 of those lines. In short, switch resources cannot be shared (non-concentrated) for the average Internet customer. This increases cost all the way through the switch, because all the concentrators and networks must be increased to handle this traffic.

When an LEC is serving an ISP in a remote, host / remote link blockages will result in slow dial tone, creating more problems for the outside plant. There are only a certain amount of cable pairs out to that remote location, and when their capacity is exhausted, the LEC must put in pair-gain devices or remote devices. If there are many ISPs, those pair-gain devices suddenly become one-to-one direct service. This, of course, presents another cost issue.

One of the hardest problems has been identifying ISPs. When the Internet first came into common use, anybody could start up an ISP, and nobody really knew where they were. LECs would have to use traffic studies to identify congestion and then try to figure out who was causing it. When they found out who and where the ISP was, the next step was to move them onto direct service from the central office or to find another solution to get them away from the remotes and the pair gains.

EAS Network Impact

The next problem is that once the ISP gets to a switch, it causes a tremendous amount of pressure on the EAS trunking network. Congestion occurs between all local offices and the ISP serving switch. Large ISPs will concentrate and buy service out of one central office. In Las Vegas, for example, there are almost a million subscribers trying to use 350 America Online lines in one switch. Once an ISP is blocking up to about 10% of the lines, then retrials begin to occur over and over. But if the LEC can keep blockage below 10%, it can handle traffic fairly well. Once the 10% level is reached, holding times increase, and every customer is competing for the same trunks.

Also, Internet usage has been mostly in the evenings. Customers call up at 6:00 p.m., and they will be on the Internet all night long, until 10 p.m. or 11 p.m. In essence, customers have changed their busy hours to evening hours, which forces the LEC to make unplanned expenditures for trunk additions. They must increase their networks over a daytime busy hour, which used to carry all its traffic. All those increases are directly attributable to the Internet.

TABLE 1

Actual CCS/Line Impact

OFFICE	TOTAL LINES	CCS/AL YE94	CCS/AL YE95	1994/1995 % GWTH	CCS/AL YE96	1995/1996 % GWTH	CCS/AL 97	1996/1997 % GWTH
A	64K	2.30	2.30	0%	2.45	4.3%	NA	NA
B	54K	3.41	3.43	0.5 %	3.50	2.0 %	3.62	3.4 %
C	46K	3.66	3.67	0.3 %	3.70	0.8 %	3.73	0.8 %
D	45K	3.08	3.24	5.2 %	3.36	3.7 %	3.59	6.8 %
E	43K	3.22	3.32	3.1 %	3.42	3.0 %	3.59	5.0 %
F	30K	3.35	3.36	0.3%	3.41	1.5%	NA	NA
G	27K	3.07	3.38	10.1%	3.72	10.1%	NA	NA
H	20K	2.90	2.98	1.3 %	3.02	1.3 %	3.07	1.7 %
J	20K	3.57	3.60	0.8%	3.68	2.2%	3.79	3.0%
K	19K	2.49	2.54	3.0 %	2.59	2.0 %	2.68	3.5 %
Average CCS/AL Growth				2.46%		3.29%		4.03%

TABLE 2

Actual Trunk Impact

OFFICE	1994 WKG TRKS	1993/ 1994 %GWTH	1995 WKG TRKS	1994/ 1995 %GWTH	1996 WKG TRKS	1995/ 1996 %GWTH	1997 WKG TRKS	1996/ 1997 %GWTH
B	4162	3.25 %	4306	3.46 %	4786	11.21 %	5746	20.10 %
C	2384	3.72 %	2553	7.08 %	2630	3.02 %	3630	38.02 %
D	1836	3.11 %	1908	3.92 %	2216	16.14 %	2648	19.50 %
E	1403	3.30 %	1451	3.42 %	1475	1.66 %	1547	4.88 %
H	408	2.91 %	454	11.27 %	550	21.15 %	742	34.91%
J	978	2.10 %	1002	2.45 %	975	-2.47 %	975	0.00 %
Average Growth		3.07%		5.27%		8.45%		19.57%

Actual Impact Based on Traffic Study Data

Table 1 shows actual CCS/access line impact data, which Sprint gathered from a number of offices in Florida and North Carolina. The study was controlled to eliminate offices in territories where a large ISP does not distort the data. For many years, these offices used to use about 2.4 100-call seconds per customer. This made planning simple: the number of customers multiplied by 2.4 determined the size of the switch. That rounds out to approximately a six-to-one concentration ratio throughout the network. The LEC could serve all its customers at once, but it only had to provide one path for every six customers.

However, this CCS/AL figure has been increasing over the last four or five years. For 1994/1995, these offices were running at 2.46 percent, a small increase. But for 1995/1996 this number increased to 3.29 percent, and for 1996/1997 the figure was up to 4.03 percent. This means that traffic is increasing at a rate of 4 percent apart from normal customer growth. If this is an LEC with a 5 percent annual growth rate, it has to deal with another 4 percent growth which has nothing to do with access-line gain. This means the LEC must come close to doubling the amount of switching equipment it would need just to add access lines to accommodate normal growth.

The actual trunk impact at six of these same offices increased from 3.07 percent in 1993/1994 to 5.27 percent in 1994/1995 (see *Table 2*). From 1995 to 1996, there was an 8.45 percent gain, and from 1996 to 1997 there was almost a 20 percent gain. These are not CCS per access lines; they are total trunks working. Only 5 percent of the gain in access lines is due to new subscribers. The other 15 percent is mostly due to the Internet.

To reiterate, the CCS per access line ratio remained constant at about 2.4 for many years, although it almost doubled in the last three years. This means that concentration ratios of 6:1 will no longer work; new ratios of 3:1 or less must be used for provisioning. As CCS per access line increases, the cost per access line must be increased when budgeting dollars for access line growth. This means that a LEC could be spending twice as much money to provide the same service. Although local trunking networks used to grow at about the same percentage rate as access lines, they will now grow between 15 percent to 20 percent, which is two to three times the rate of growth in access lines.

Summary

Internet service and other factors are causing a 4 percent per year increase in the use of switching equipment, trunking, and facilities. Provisioning costs for a 4 percent increase in CCS per access line is about the same as provisioning for an 8 percent increase in access-line growth. A LEC that is budgeting equipment for an 8 percent increase in access-line growth will have to add at least another half of that just for Internet growth. On the EAS trunking side, growth is being forecast at 15 percent–20 percent, where traditionally it has been in the 8 percent to 10 percent range for these same offices.

The Internet is just the beginning of increased usage. There are other factors that will cause even more headaches. Many more people are working at home, which means they are on the telephone or the Internet for 8 to 10 hours per day or more. Web TV will have a major impact if it truly gets going. Also, competitive local exchange carriers (CLECs) require an increase in the trunking network because of local number portability (LNP). In order to connect to a CLEC as a result of LNP, an LEC may well have to switch a given call three times between three different switches. Local number portability is a whole new ball game, and it will cost everybody a lot more money than they think. Finally, educational services via the Internet will be another major impact on network usage.

In the future, one of the LEC's key needs will be to provision switches and networks to meet engineering objectives. Whether the objective is one blocked call in a hundred or one blocked call in a thousand, it still must be met, because as soon as an overflow situation occurs, everything goes out the window.

Also, LECs should work with ISPs to provide connectivity through ISDN–PRI or DS-1 access in order to reduce congestion. They should be taken off old line-ended modems and moved to something a little faster in the main switches. Also, ISPs should be encouraged to collocate their systems physically inside the LEC's central office or very close to it in order to save outside plant costs.

Determinants of Success of Telemedicine Programming

Joel L. Nitzkin, M.D., D.P.A.
Larry R. Vinson, Ph.D.
Naisu Zhu, M.D.
Jayaraman Rao, M.D.
David Diloreto, M.D.

Introduction and Historical Perspective

The term "telemedicine," as used in this paper, refers to physician consultation to patients at remote sites, using live, interactive video communications.

Telemedicine is now about 40 years old—dating from a mental health project in Nebraska in 1958. Of all of the projects begun in the 50's, 60's, 70's, and early 80's, only one still survives—on an Indian reservation in New Mexico. There has been a veritable explosion in telemedicine during the 1990's—spectacular growth fueled by improvements in technology, reductions in cost, and the dynamics of the medical marketplace. Despite this growth, many telemedical stations still average less than one clinical encounter per week. Some have none for months at a time. Except for projects within the American military and the Texas prison system, none process enough clinical encounters to be cost-effective on a direct-service basis.

Administrative and interpersonal factors are the major determinants of success of telemedicine programming. Patients are uniformly enthusiastic about telemedicine. Physicians and administrators remain skeptical. Physician skepticism relates to the perceived quality of care. Administrative skepticism is based on financial considerations: telemedicine systems can threaten the viability of the medical practices of the specialists at regional centers that patients would otherwise attend. Local residents gave lackluster performance ratings to their local health care delivery system. Residents of Lake Charles did not think a telemedicine program would improve the local health care delivery system, but rural residents and charity patients thought telemedicine would significantly improve quality and access to care. A demonstration of the telemedicine system convinced charity patients that telemedicine would be even more beneficial than they had thought. Physicians and administrators were much less accepting of telemedicine than urban or rural residents, charity clinic patients, and directors of nursing.

This paper will explore the reasons for the low utilization rates from patient, physician, administrative, and community perspectives. It will also offer suggestions to vendors of telemedicine systems to help enhance the success of telemedicine programming.

Description of the LSU TELEMEDicine Research Project

The Louisiana TELEMEDicine Research Project® was conducted in two phases. Phase I, from February through April of 1994, consisted of a preliminary research project in which fully configured telemedicine stations were set up in adjacent rooms at University Hospital. Physical therapists and physicians, by specialty, were invited to bring groups of patients, examine them telemedically and conventionally, and compare the results. Phase I utilized the then current model of the CLI/Rembrandt II/VP codec, with multiple attachments and the CTX-Plus proprietary algorithm. Transmission between the two codecs was at full T-1. The final Phase I data set consisted of data on 58 patients and 33 physicians and therapists, representing physical therapy and five medical specialties.

Phase II consisted of trial clinical encounters between University Hospital in New Orleans, Louisiana, W.O. Moss Regional Medical Center in Lake Charles, Louisiana, and South Cameron Memorial Hospital in Cameron, Louisiana. Phase II Utilized VTEL codecs, with attachments similar to Phase I, also at Full T-1. Phase II began in August of 1994 and ran through August of 1995. While about 100 clinical encounters took place in Phase II, patient and physician satisfaction data were only gathered at the initial encounter each calendar quarter. Since some patients were seen multiple times on the system, the satisfaction database reported herein is limited to 53 patient satisfaction forms and 32 physician satisfaction forms.

Routine clinical use began in September of 1995, and was ongoing at the time of publication. The project switched back to CLI codecs. Since no satisfaction data are now being gathered, none are reported herein.

Patient Perspective

Patient satisfaction data are shown in *Tables 1, 2* and *3*. In both Phases I and II, patients were enthusiastic about telemedicine. This is remarkable for the Phase I patients since, as research subjects, they were not paid to participate, and they received no medical benefit from these encounters. Each patient had to undergo multiple examinations by multiple physicians during the two-hour clinic period. Many felt honored to have the

TABLE 1

Patient Satisfaction—Phase I—35 Ophthalmology Patients

Variable	No. Positive	Percent Positive
Confident doctor understood problem	34	97.1%
Could say or ask anything I wanted	34	97.1%
Prefer Telem Consult to two-hour drive	31	88.6%
Overall Feelings	34	97.1%

TABLE 2

Overall Patient Satisfaction—Phase II

Ranking	Number of Patients	Percent of Patients
Excellent	50	94.3%
Good	2	3.8%
Average	0	
Fair	1	1.9%
Poor	0	
Totals	53	100.0%

TABLE 3

Patient Comments—Phase II

```
PT   BSITE DATE     OVRCOM
----- ---  -------- ------------------------------------------------------------

SC004  UH  08/25/94 first hesitant, now impressed, thankful for dr sanders
SC006 WOM  09/01/94 other than too many in w/ you during exam I enjoyed it.
SC015  UH  10/14/94 It was great. It saved a lot of time and money.
SC016  UH  11/03/94 Telemed is a very good idea.
SC017  UH  11/09/94 For a serious illness this should not be used.
SC018  UH  11/10/94 Telemed very positive.
SC021  UH  11/22/94 Very good idea. Hope you keep it.
WOM001 UH  09/09/94 very nice and conveniently located
WOM002 UH  08/30/94 prog&access super,pts can't travel,blends trad&high tech.
WOM003 UH  08/30/94 from what I've seen,i'm very positive about it
WOM004 UH  08/30/94 good for cases who don't need physical exam
WOM005 UH  08/30/94 trouble understanding the dr. but I think I have a bad ear
WOM006 UH  09/09/94 I was very satis w/ visit & it is a great medical advancement.
WOM007 UH  09/23/94 it should be continued
WOM009 UH  09/23/94 great
WOM011 UH  09/23/94 I like telemed. It saved me a whole day travel to n. Orleans.
WOM012 UH  09/23/94 I like telemed! I think it will be very beneficial to all.
WOM014 UH  09/23/94 was very satisfied with my visit
WOM015 UH  10/17/94 grateful to talk to dr.rao;to get opinion and possible tx.
WOM016 UH  10/17/94 I was treated very well by the staff at the hospital.
WOM020 UH  10/28/94 It's progress for medicine at a welcomed time.
WOM021 UH  11/28/94 Very pleased;worthwhile project for SWLa pts. Thank you
WOM022 UH  11/28/94 I think this is wonderful.
WOM025 UH  12/12/94 makes quality care within reach for many
WOM027 UH  12/12/94 very well pleased with tm;look forward to meeting dr.rao
WOM035 UH  02/15/95 saves time & discomfort of travel;like rapport with doctor.
WOM036 UH  02/15/95 I hope it is here for good.
WOM037 UH  02/15/95 very pleased with program
WOM040 UH  02/15/95 excellent, superb!
```

opportunity to participate in this research. To these patients, the benefits of telemedicine were apparently obvious.

The "overall satisfaction" comments of all of the Phase II patients are shown in *Table 3.* These patients did receive clinical benefit from these visits, as if they had seen the physician in person. Each patient seen at the University Hospital base site was saved a trip to New Orleans that would have taken five hours each way by car. The patient seen at W. O. Moss was saved a 40 minute trip to Lake Charles from Cameron.

The results of this study are typical of those of other telemedicine programs. Patients have felt comfortable with the telemedicine encounters and have benefited from the convenience of not having to travel to a distant city to see a physician specialist. Few patients felt that telemedical encounters were less private than conventional medical encounters, despite the presence of nursing staff in both base and remote stations, and the knowledge that the encounters were being videotaped for the medical records. Some patients actually felt more comfortable with the physician at a distant site.

For some patients, such as those with Parkinson's Disease and other movement disorders, five hours in a car is a very trying experience. The difficulty of the trip severely inhibits access to care. For these patients, telemedicine proved very beneficial and enabled the physician to see them more frequently than would otherwise have been feasible. For patients coming to the medical center for highly specialized surgical procedures, the telemedicine program reduced the need for long-distance travel for preoperative and postoperative clinic visits.

Physician Perspective

The data on physician satisfaction is shown in *Tables 4, 5* and *6.* These data indicate an opinion about telemedicine that is positive in tone, but not as enthusiastic as those expressed by patients. Many physicians feared that clinically significant information could be lost by not having more personal contact with the patient, and by the limitations of the audio and video components of the telemedicine system.

It is important to note that these physicians were self-selected. Their attitude was much more positive than that of colleagues who declined the invitation to participate in this program. Recognizing in advance that a lack of health insurance reimbursement for telemedical encounters might discourage some physicians from participating, a small sum of money was secured, from which each of the participating physicians was reimbursed at the then current Medicaid rate. Thus, lack of reimbursement was not a factor limiting physician satisfaction in this study.

Researchers found throughout their experience, much of which consisted of informal interactions with both participating and nonparticipating physicians, that physician satisfaction varied with the relationship of the physician to the telemedical program.

For specialists, the most substantial issue was fear of the loss of clinically significant information. This was due to the lack of more personal contact with the patient and uncertainty as to the audio and video limitations of the system. This, physicians feared, could significantly increase the likelihood of a diagnostic or therapeutic error. The strange "feel" of the remote encounter and the need to rely on distant staff to position the patient increased their feelings of uneasiness with telemedicine. This in turn led them to worry about the lack of personal contact reducing the strength of the physician-patient relationship.

Yet another factor for these specialists was skepticism about the motives of the proponents of the telemedicine programming. Was this being used as a means to boost medical revenues to the facility without regard for quality of care? Would this process, while benefiting the facility, expose the physician to malpractice litigation?

The duration of the clinical encounter and physician productivity were also issues. Medical histories took much less time telemedically, as the system seemed to discourage social "small talk" between patient and physician. At the same time, at least for some specialties, physical examination took much longer. Thus, for some medical specialties, many more patients could be seen telemedically in a one-to-two hour telemedical clinic while for other specialties, fewer patients could be seen in the same period of time.

For all participating physicians, reliability of the telemedical equipment, convenience of access to the telemedicine suite, and quality of staff support were critical. Nothing was more frustrating for physician or patient than to have to cancel the encounter because of an equipment-related problem. While equipment-related glitches occurred on an almost daily basis, the researchers' staff was able to deal with most of them very quickly, so medical visits did not have to be rescheduled.

When designing the program, researchers anticipated that most telemedical encounters would occur in response to referral by the local rural general practitioners. These physicians, while favorable to the telemedicine program, rarely felt the need for such referral. Most Phase II encounters were initiated by the specialists to reduce the need for travel by patients already under their care.

Some rural general practitioners expressed the concern that telemedicine, once well developed, would be used by facility administrators to displace local physicians by nurse practitioners backed up by specialist physicians from remote medical centers. Given stories of physician unemployment in urban areas due to managed care, and the aggressive manner in which some health care administrators were pursuing cost reduction at physician expense, these fears were very real.

The physicians most threatened by telemedicine were specialists practicing in the studied rural areas and regional centers. They saw themselves being put out of business by the "big city" medical centers coming to town telemedically and selling their services cut-rate. Some of these physicians are leaders in local and state medical societies and are doing everything

TABLE 4
Physician Satisfaction—Phase I—Ophthalmology

Variable	Responses	Number Positive	Percent Positive
interest in serving as telemedical consulting physician	21	19	90.5%
interest in serving as telemedical referring physician	21	16	76.2%
as a patient, I would be willing to accept diagnosis and treatment based on a telemedical examination	22	6	27.3%
I am comfortable working with telemedical equipment	22	7	31.8%
My overall rating of telemedicine is	20	16	80.0%
Average	21.2		60.4%

TABLE 5
Physician Satisfaction—Phase II

Ranking	Number of Physicians	Percent of Physicians
Excellent	17	53.1%
Good	12	37.5%
Average	1	3.1%
Fair	2	6.3%
Poor	0	
Totals	32	100.0%

TABLE 6
Physician Comments—Phase II

```
SITE DATE     TMDNU  OVERCOM
---- -------- ------ ------------------------------------------------------------
UH   08/01/94 SC002  more timely consult and improved f/u for remote patients.
UH   08/02/94 SC003  need more exposure
UH   08/25/94 SC004  great for neurology
UH   08/30/94 WOM002 can't get used to it, not a major problem
UH   09/29/94 SC014  Great potential but may have limitations in ophthalmology
UH   10/14/94 SC015  interesting experience! Hope to use again for f/u patients.
UH   10/28/94 WOM019 ok for pt/md discussion--not phys exam; video poor
UH   11/03/94 SC016  good for pt/family, not for md; overall telemed is good
UH   11/30/94 WOM023 2.for 1st evals;good for f/u;new pt:direct contact
UH   01/05/95 SC027  program is valuable in providing comprehensive care.
WOM  09/20/94 SC011  saved pt a trip to mrmc surg clin for diagnosis
WOM  09/23/94 SC012  yes, help pt control diabetes
WOM  10/28/94 WOM019 pts would benefit;feel more relaxed;chance to see md's
```

within their power to keep both managed care and telemedicine from their respective communities.

Administrative Perspectives

There are at least four distinctive administrative perspectives from which to view telemedicine. Each has a unique set of direct and indirect costs and benefits. From each of these points of view, telemedicine has the potential to be beneficial, but only if the system is planned to secure these benefits, and the utilization rate is high enough to justify equipment, transmission, and staffing costs.

The first of the four perspectives is that of the rural hospital or other remote (patient) site. If, by offering telemedical access to specialty physicians, the rural site can increase community confidence in its ability to handle more complex medical problems, more of the medical dollars can be retained locally by generating greater inpatient and outpatient revenues. The benefit to be secured, from a facility perspective, is almost entirely indirect.

The second of the four perspectives is that of the urban medical center. If, by offering telemedical services, it can increase its referral base, then it can fill more of its hospital beds in this highly competitive and over-bedded medical marketplace. Here too, the major benefit is indirect. If the medical center can keep an average of one more bed filled each day, year-round because of its $100,000 per year telemedicine program, it will generate about $365,000 in revenues. This is a far greater sum than the medical center could secure through insurance reimbursement of the telemedical encounters.

The third of the four perspectives is that of the HMO or other managed care plan. Here the plan is both payer and provider. The major benefit of telemedicine is reduction of patient-related travel. Since these costs are not usually covered by health insurance, this is not a benefit to the provider or payer. Telemedicine's major benefits to the HMO are also indirect. They relate to the degree to which a telemedicine program will enhance the HMO or managed care plan's marketing in rural areas.

The fourth perspective is that of Medicare, the largest payer for health care in the United States. Medicare has not routinely funded telemedical encounters to date because they fear paying for large amounts of equipment that would sit unused. Medicare is also concerned that by funding telemedicine programming, they might be subsidizing substandard medical care. No definitive research has yet been conducted to demonstrate that the quality of telemedical diagnosis and treatment are fully equivalent to conventional medical practice. Current research, however, suggests that the quality of telemedical care can equal that of conventional medical care, if the staff have been well trained and the participating physicians are fully aware of the limitations of the equipment in use. These are conditions that may or may not be met by any given telemedical program, due in part to the fact that formal guidelines have not yet been developed for non-military use.

Yet another issue, from the Medicare perspective, relates to the enhancement of access to care. This enhanced access is the primary benefit of telemedical programming. Medicare sees the possibility that telemedicine could dramatically increase medical utilization by rural clients without corresponding enhancements of health status or reductions in hospital care. This in turn could significantly increase cost without apparent benefit, at least from the perspective of this payer.

It appears that telemedicine will cost rather than save money. Telemedicine can reduce costs only where transportation is extremely costly, and the cost of the transportation is picked up by the health care provider. This limits the possibility of cost savings to the military, prisons and jails, ships at sea, and other hard to reach populations. Even there, however, the volume must be high enough to overcome the fixed costs of the telemedicine system.

Community Perspective

The community perspective regarding telemedicine programming varies by socioeconomic status and location. *Tables 7, 8* and *9* show community survey data gathered by the LSU program. The "urban" and "rural" data were gathered by telephone interview of residents of Lake Charles, Louisiana, and surrounding rural parishes (counties). These respondents were presumed to be middle class.

The "clinic" patients were those being seen on the day of the survey at W. O. Moss Regional Medical Center, the "charity" hospital serving this region of Louisiana. Some were urban, some were rural, and all were poor.

The data in *Table 7* show that all three groups of patients, "urban," "rural," and "clinic," had low opinions about how well their local and regional health care delivery systems were meeting their medical service needs. On a quality of performance scale from zero to four, almost all parameters for all groups averaged between two and three. These ratings are clearly lackluster, and point out the importance of local hospitals and medical groups taking steps to enhance the confidence of the local community in the quality of service being provided.

The data in *Table 8* address the extent to which implementation of a telemedicine program would improve the local health care delivery system, from the perspective of the three segments of the community studied. Nine dimensions of the health care delivery system were included in the study questionnaire. The "urban" group did not believe that telemedicine would have a major favorable impact on any of the dimensions. With respect to the "caring" quality of the care, the "urban" subjects thought telemedicine would make the current situation worse. Both the "rural" and "clinic" groups were much more positive, noting major improvements for six and seven of the nine parameters, respectively. The "clinic" group seemed to feel that anything would help, while the rural group seemed to feel that the benefits of a telemedicine program were self-evident.

Table 9 shows the difference in perception of the "clinic" group before and after a demonstration of the telemedical

TABLE 7

(Lake Charles Community Survey) Importance and Grade of Characteristics of Health Care Delivery System

Characteristic	Importance (to survey respondent)			Grade (how well local system performs)		
	urban (N=274)	rural (N=168)	clinic (N=66)	urban (N=274)	rural (N=168)	clinic (N=66)
caring manner	4.71	4.83	**4.88**	2.17	2.04	2.23
see doctor when needed	4.78	**4.86**	4.80	2.14	2.15	2.45
short wait in wtg rm	*4.46*	*4.48*	*4.53*	2.73	2.57	**3.25**
receive care locally	4.73	4.82	4.83	2.03	2.47	2.23
receive top quality care	**4.84**	**4.85**	**4.91**	2.18	2.38	2.11
able to afford care	4.81	**4.84**	**4.86**	2.56	2.51	2.54
privacy/confidentiality	**4.84**	4.83	4.68	*1.66*	*1.64*	*1.58*

Scale: Importance -- 0-5; Grade -- 0-4

TABLE 8

(Lake Charles Community Survey) Anticipated Impact of Telemedicine on Local Medical Care System—Average Score

Characteristic of Health Care Delivery System	urban (N=274)	rural (N=168)	clinic (N=66)
Access to family doctor	3.74	**4.15**	**4.26**
Access to medical specialist	3.78	**4.35**	**4.48**
Timely access to doctors	3.66	**4.09**	**4.41**
Access to doctors without having to travel	3.98	**4.51**	**4.71**
Affordability of medical care	3.33	3.58	3.68
Personal, caring medical care	*2.96*	3.31	3.88
Quality of medical care available locally	3.58	**4.14**	**4.35**
Confidentiality	3.02	3.46	**4.08**
Access to health related education	3.52	**4.04**	**4.39**
Percent of Parameters with Major Favorable Impact	0.0%	66.7%	77.7%
Percent of Parameters with Unfavorable Impact	11.1%	0.0%	0.0%

TABLE 9

W. O. Moss Pretest-Posttest Results

"Clinic" Group; N = 66

Item	Pretest	Posttest	p <
Access to family doctor	4.15	4.77	.0001
Access to medical specialists	4.43	4.87	.001
Timely access to doctors	4.34	4.75	.002
Access to doctors without having to travel	**4.63**	**4.81**	**.132**
Affordability of medical care	3.50	4.20	.0001
Personal, caring medical care	3.79	4.34	.002
Overall quality of medical care available locally	4.30	4.70	.003
Confidentiality	3.98	4.38	.008
Access to health related education	4.32	4.70	.003

equipment. The scores improved for all nine parameters. Most improvements were highly statistically significant except the one relating to travel, which scored very high on the initial questionnaire. These findings, in turn, suggest that marketing of the telemedicine program within the community might improve the image of the sponsoring facility.

Table 10 summarizes some of the data from the three community groups and compares them with data from very small samples of physicians, administrators, Directors of Nursing within the Lake Charles region, and Directors of Nursing from other LHCA ("charity") hospitals. There are dramatic differences in perception between the community samples and Directors of Nursing on one side, and the physicians and administrators on the other. This difference is not a trivial one, as the administrators decide whether a facility is to have a telemedicine program and how the program is to be managed.

The physicians did not share in the opinion that telemedicine would improve quality of care, as opposed to 82 percent to 98 percent of the community groups, who felt otherwise. The physicians and the administrators felt that telemedicine would not improve access to care. Eighty-seven percent to 96 percent of the community groups and 75 percent and 88 percent of the two groups of Directors of Nursing, however, felt that it would improve access. Major differences were also seen in response to questions regarding reasons that locals seek care elsewhere and their preferences for telemedicine over travel to a distant site for medical care. The most positive responses from the physicians and administrators were for the favorable impact of telemedicine on continuing medical education. Even here, however, the responses were lackluster, running only 54 percent for the physicians, and 50 percent for the administrators.

Steps that Should Be Taken by Manufacturers, Integrators, and Vendors to Enhance the Success of Telemedicine Programming

The recommendations that follow reflect the personal opinion of this research's principal investigator (JLN). They are based on experience within Louisiana, and on informal discussions with the staff of other telemedicine programs.

1. Vendors should consider the following:

   - Administrative and interpersonal factors are the major determinants of success of telemedicine programming, rather than the technical sophistication of the equipment.
   - There is a dynamic interplay between the quality and sophistication of the telemedicine equipment and the costs and benefits of the telemedicine program. The better the equipment, the lower the maintenance, training, and staffing costs. Increasing the sophistication of the equipment should shorten the time required to conduct the physical examination portion of the medical encounter.
   - It is unlikely that any nonmilitary, non-correctional telemedicine program will ever be cost-effective based on the number of clinical encounters processed. The benefits will be almost entirely indirect.
   - Rural and urban facilities, physicians, administrators, and payers each have different perspectives of telemedicine programming. What is a benefit to one is often seen as a cost to another.

2. Vendors should orient marketing efforts to the concerns of the administrators and physicians, from their respective

TABLE 10

Summary of Louisiana Survey Data—Selected Data Points

Item	Clients/Patients			Lake Charles Region			LHCA
	urban (N=274)	rural (N=168)	clinic (N=66)	MD N=11	Admin N=6	DON N=4	DON N=16
telemed would improve quality of care	82%	95%	98%	no	(not asked)		
telemed would improve access to care	87%	96%	95%	18%	33%	75%	88%
Reasons Locals Seek Care Elsewhere:							
lack of confidence in local providers	critical issue	mid-range response		54%	17%	50%	25%
physicians overbooked	minor problem	not a problem		54%	0%	25%	6%
service not available locally	major problem	mid-range response		9%	17%	(not asked)	50%
Preferences							
patient would prefer telemed over trav.	(These questions were not asked in community surveys)			18%	33%	75%	88%
respondent prefer telemed over trav.				27%	33%	50%	75%
favorable impact on contin. med. ed.				**54%**	**50%**	**75%**	**81%**

points of view, with emphasis on community support and indirect benefits.

3. Vendors should urge administrators to involve physicians in the early planning stages of the telemedicine program and allow them a major role in the selection of equipment. Failure to do so may result in physician opposition to the program.

4. Since physician referrals for specialty telemedical consultations will not generate enough telemedical encounters to justify program costs, program plans should include continuing professional education for physician and non-physician staff, and partnership arrangements within which medical center specialists conduct telemedicine outpatient clinics at remote facilities.

5. Vendors should improve the reliability of the telemedical systems to reduce staffing and maintenance costs and improve physician willingness to depend on the system. The goal should be for the entire telemedicine system to equal the reliability and consistency of performance of home television sets. Failing this possibility, vendors should advise potential buyers of necessary staffing patterns and maintenance costs based on experience at other sites.

6. Telemedicine research and training programs should be developed at industry expense. This would have two benefits. The first would be an improvement in the quality and consistency of the medical care provided. The second would be a more in-depth orientation to the systems for both administrators and physicians, which should result in more favorable opinions and more support for the telemedicine program. Elements to be considered in these research and training programs include the following:

   - Through partnership with a medical center, each component should be on a wide range of acutely ill patients by the widest possible range of medical specialists.

   - Find out how the examination is to be done if the telemedical examination is to be the equivalent of a conventional examination.

   - Determine how data are best captured for the medical record as still image, moving image, or sound.

   - Training protocols are required for:
     - remote (patient site) nurse
     - base (consulting physician) nurse
     - consulting physician

   - NOTE: The telemedicine nursing staff should also be trained to troubleshoot technical problems with the telemedicine system and address the frequent minor glitches as they occur. Since the nurse will usually be the only on-hand assistant to the physician, he or she might be the only person capable of saving an encounter, should a problem occur.

   - Training protocols must be developed for each component and equipment configuration.

   - Each base and remote site must be provided with means and protocols for future staff training—most likely in loop-back mode—so patients can be seen both conventionally and telemedically by each physician and nurse in a single training session.

   - There should be better standardization of codecs and components, with emphasis on the need to transmit diagnostic quality audio and video between codecs from different manufacturers.

8. More information should be provided on system limitations and needs for staff support within the marketing programs, especially regarding the following points:

   - accuracy of physical examination findings

   - protocols provided for patient examination

   - data provided on comparative durations of physical examination for conventional and telemedical examination

   - process for capturing images and sounds in medical records

   - reliability of overall system as to
     frequency of system malfunctions;
     staffing and staff training required to secure optimal reliability of system operations; and
     availability of technical support for problems that may be equipment or transmission related, or relate to the interface between the equipment and transmission lines

9. Vendors should document utility of telemedical equipment for training and quality assurance for on-site medical care, and market systems for this purpose

Annotated Bibliography

Telemedicine, General

Allen, A. and D. Allen. "Telemedicine Programs: 2nd Annual Review Reveals Doubling of Programs in a Year, " *Telemedicine Today*, no. 1 (spring 1995): 10-18. Most telemedicine systems only see one to five patients per month; nowhere near enough to be cost-effective. The problem is usually stated in terms of physician reluctance to use the system.

Perednia, D. A. and A. Allan. "Telemedicine Technology and Clinical Applications," *JAMA* 273, no.6 (February 8, 1995): 483-488. A general introduction to telemedicine. (Full text of AMA journal papers is available on-line at http: / / www.ama-assn.org)

Sanders, J. H. and R. L. Bashshur. "Challenges to the Implementation of Telemedicine," *Telemedicine Journal 1*, no. 2 (summer 1995): 115-124. A discussion of problems and barriers relative to implementation of telemedicine programming.

Web site to access telemedicine-related literature:
http: / / www.telemed.org or http: / / www.tie.telemed.org—The Telemedicine Information Exchange run by the Telemedicine Research Center (Dr. Perednia) , Oregon Health Sciences University, Portland Oregon.

History of Telemedicine

Bashshur, R. L. and P. A. Armstrong. “Telemedicine Journal: A New Mode for the Delivery of Health Care,” Inquiry 13, no. 3 (September 1976): 233-244. An excellent and in-depth review of the history of telemedicine to 1976, with in-depth discussion of the factors predicting the success and failure of telemedicine projects. This paper is remarkable for the similarity of the issues and problems faced by today's telemedicine projects. An extensive set of references is provided.

Brauer, G. W. “Telehealth: The Delayed Revolution in Health Care,” *Med Prog Technol.* 18 (1992): 151-163. A much more recent historical review of telemedicine.

Wittson, C. L., D. C. Affleck, and V. Johnson. *The Use of Two-way Television in Group Therapy*, Omaha: Nebraska Psychiatric Institute, University of Nebraska College of Medicine, 1961. 1959 microwave linkage used for telepsychiatry consultations between the Nebraska Psychiatric Institute in Omaha and the state mental hospital 112 miles away.

Technology and Health Care Cost—with Emphasis on Telemedicine

Grigsby, J. “Lies and Damned Lies,” *Telemedicine Today* 5, no. 1 (January / February 1997): 34-35. A guest editorial by an expert in health policy research who has worked in the telemedicine field for many years blasting claims that telemedicine systems are cost-effective. Author can be reached at grigsby@essex.hsc.colorado.edu.

Nitzkin, J. L. “Technology and Health Care—Driving Costs Up, Not Down,” *IEEE Technology and Society Magazine* 15, no. 3 (fall 1966): 40-45. This paper explodes the myth that investment in technology of any sort, but most particularly telemedicine and electronic medical records, is likely to reduce overall health care costs, and explains why. Extensive list of references.

Reliability of Telemedicine Examination

Nitzkin, J. L., N. Zhu, and R. L. Marier. “Reliability of Telemedicine Examination,” *Telemedicine Journal* 3, no. 2 (summer 1997): 141-157. Telemedical examination, across a broad range of medical specialties, can be as reliable as conventional examination if the participating clinicians (doctors, nurses, therapists) have had experience or training in telemedicine and are familiar with the limitations of the telemedicine system in use. Without such experience or knowledge, they could miss physical findings of major clinical importance. These findings raise doubts about the reliability of occasional telemedicine consultations by clinicians inexperienced in telemedicine. Recommendations are made for training and research, and an extensive bibliography is appended.

Intelligent Agents and Their Future on the Internet

Rajesh C. Phoha
Professor of Computer Science
Jawaharlal Nehru University

Vir V. Phoha
Assistant Professor of Computer Science
Northeastern State University

Abstract

Rapid expansion of the World Wide Web, its dynamic, diverse, and heterogeneous information sources, and its accessibility to a large number of nontechnical users has created an intense need for software tools that are easy to use, can search for and present information in a format and combination that best fits user needs, and can perform repetitive and mundane tasks. To meet this need, the past few years have seen the emergence of software components, "intelligent agents," that perceive, reason, act, and are autonomous, sociable, and cooperative. This paper presents the concept of a software-based intelligent Internet agent and reviews current research and several Internet agents that are commercially available. The future of intelligent Internet agents is also addressed; the paper concludes with a discussion of desirable features for Internet agents.

Introduction

Rapid increase in the use of World Wide Web and its availability to novice computer users combined with an explosion of available information on the Internet has led to an increasing need for simple and easy-to-use tools for searching for information and the automation of routine jobs. The World Wide Web is an open system with diverse information sources, which have different mechanisms for storing and presenting information. The information and sources of information on the Web may appear or disappear arbitrarily. Further, the vast amount of information on the Web needs to be advertised, found, collated and presented, and updated. To develop a software system that will not only manage this information but also meet the users' needs by collecting and presenting the information in a desirable form presents a formidable task.

Recently there has been a spurt of activity, both in university computer science departments and in the industry, in developing intelligent agents that can serve as surrogate users. Despite the absence of a unified theory to form the basis of intelligent agents, the field is rich. This apparent paradox arises from the fact that researchers in this field are able to draw from well-established areas in computer science such as artificial intelligence, design of programming languages, computer communication, and from other allied fields such as psychology and sociology.

It is natural then to ask, "Where are we and where should we be heading?" This paper briefly reviews (1) what an agent is and how an agent differs from a program, (2) some currently available Internet agents, and (3) some desirable features and the future of Internet agents.

What Is an Agent?

Researchers have defined the agent in a variety of ways. However, there is a universal agreement that agents are software components with certain essential characteristics. Some commonly accepted characteristics of intelligent agents are perception, autonomy, communication, and reasoning. In the following paragraphs, some commonly accepted definitions of software agents are presented. (For a comparison of definitions of "intelligent agent," see "Is It an Agent or just a Program: A Taxonomy for Autonomous Agents" by Franklin and Graesser or Petrie's "Agent-Based Engineering, the Web, and Intelligence."[1,2])

According to Russel and Norvig, "An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors."[3] Marc Belgrave describes agents as nondeterministic automatons that can employ attributes such as emotion, belief, and commitment.[4] Some agents can even emulate real people. Barbara Hayes-Roth writes,[5] "Intelligent agents continuously perform three functions: perception of dynamic conditions in the environment; action to affect conditions in the environment; and reasoning to interpret perceptions, solve problems, draw inferences, and determine actions." The agent from IBM's intelligent agent strategy white paper defines agents as software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user's goals or desires.

Some key characteristics of intelligent agents that differentiate them from other types of applications are:

- *autonomy* (agents should be able to operate—make decisions—on their own and have some kind of control over their actions and internal state)
- *proactiveness* (an agent should be able to exhibit goal-directed behavior by taking initiative)

- *communication ability* (an agent must be able to communicate with its environment)
- *reasoning* (an agent should have the ability to infer and extrapolate based on current knowledge and experiences in a rational and reproducible way)
- *adaptive behavior* (agents must be able to adjust to and examine the external environment—i.e., the World Wide Web—and the success of previous actions and adapt their actions to improve the probability of achieving goals)
- *trustworthiness* (an agent should instill in the user a degree of confidence in its action and reports)
- *personalizability* (an agent should have some component of learning and memory, and it should be educable)
- *cooperation* (the user specifies what actions to perform and the agent specifies what it can do and provides results)

Difference Between an Agent and a Program

The difference between an agent and a program, fortunately, is rather easily clarified. An autonomous agent is a system situated within and part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to affect what it senses in the future. All software agents are programs, but not all programs are agents. For example, although a typical spellchecker is simply a program, a spellchecker program that watches as a user types and corrects on the fly is an agent.

Can a script be considered an agent? A script is a structure that represents some consistent sequence of events and is generally a set of preprogrammed actions, usually to be executed in a linear, deterministic fashion. Given the same set of conditions and the same goal, a script will always produce the same result. An agent, on the other hand, may produce different results based on its current emotional state, beliefs, and so on. Agents also learn and adapt to new situations in a manner that clever scripts can only approximate.

Internet Agents

Today's Web is dominated by applications such as information access, information filtering, electronics commerce, and workflow management. These applications have a need to manage open, fast-changing, and distributed information.

Intelligent agents solely dedicated to applications on the Internet may be classified as Internet agents. An Internet agent may be further characterized by its ability to handle electronic mail, filter electronic news (e.g., Usenet Netnews), search the Internet, manage electronics commerce, and schedule meetings from heterogeneous information sources. Because of the changing nature of the Internet, the agent should be both extensible and flexible. In the following section, some common Internet agents are discussed.

Common Internet Agents

Today's agents are already sophisticated enough to shop for the user. A product called Shopbot searches a store's Web site to find the product catalog, learns the format in which product descriptions are presented, and, from these descriptions, extracts product attributes such as price.[6] Some of the more common agents are Verity Inc.'s Topic, Agentware, RadioWeb, and Highlights.[7] Topic gathers new information on users' favorite Web sites and caches it on their hard drives. Agentware has been designed to allow off-line keyword searching of the Web for later retrieval.[8] It operates using neural net and fuzzy logic techniques. RadioWeb is the first agent-based World Wide Web wireless delivery service to operate across nationwide wireless data networks.[9] It typically takes less than 15 seconds for the RadioWeb agent to pull down the remote Web page, optimize the page, and deliver it to the mobile customer. Highlights continuously monitors, retrieves, and highlights relevant updated information from Web documents.[10]

The inner workings of the Internet Softbot, designed at the University of Washington, provide a good example of how these Internet agents achieve their goals. The Softbot uses search inference and knowledge to determine how to satisfy the user's requests.[11] The Softbot can handle goals specified in an expressive subset of first order logic. Conduction, disjunction, negation, and universal quantification can be composed to specify goals for the Softbot. The Softbot's planner automatically synthesizes and executes plans to achieve the inputted goals.

Another interesting example of techniques are those used by Internet learning agents (ILA).[6] The Internet learning agent relies on a combination of test queries and domain-specific knowledge to automatically learn descriptions of Web services. The learned descriptions can then be used to enable automatic information extraction by intelligent agents.

Brian A. LaMacchia has invented a type of resource-discovery tool, designed to help users extract useful information from the Internet, called Internet Fish.[12] An Internet Fish will initiate research, continue to discover new sources of information, and keep tabs on new developments in a particular topic. Internet Fish differs from other Internet discovery systems in that it is persistent, personal, and dynamic. They incorporate deep structural knowledge of the organization and services of the Internet and are also capable of on-the-fly reconfiguration, modification, and expansion. *Table 1* lists some commercial current agent systems.[13]

The Future of Internet Agents

The future of Internet agents is linked to the future of the Internet. First, an increasing use of personal and mobile computing devices and interactive television coupled with multimedia capabilities will make voice understanding and voice communication an important part of Internet agents. Second, Internet agents will have sophisticated indirect management capabilities.[14] Indirect management occurs when the user and the agent both supplement each other and cooperate to monitor and perform various tasks instead of relying solely on a

TABLE 1

Agent Systems, Their Manufacturer, and the Manufacturer's Web Address

AGENT SYSTEM	WEB ADDRESS	MANUFACTURER
Live Agent and Search Agent	http://www.agentsoft.com	Agentsoft
Logicware	http://www.crytaliz.com/logicware/logicware.htm	Crystaliz
Passport Office, Community Navigator, Catalog Navigator	http://www.firefly.com and http://www.firefly.net	Firefly Network
Virtual IP Network Agent Applications	http://www.ftp.com/product/4agent.htm http://www.ftp.com/product/vip_wp.htm	FTP Software
Aglet Workbench	http://www.trl.ibm.co.jp/aglets	IBM
KYMA Atlantis	http://www.info.unicaen.fr/~serge/kyma.html	KYMA Software
Carnot	http://www.info.unicaern.fr/~serge/kyma.html	Microelectronics and Computer Technology Corporation (MCC)
Microsoft Agent	http://www.microsoft.com/intdev/agent	Microsoft
Concordia	http://www.meitca.com/HSL/Projects/Concordia	Mitsubishi
Oracle Mobile Agents	http://www.oracle.com/products/networking/mobile_agents	Oracle

user-initiated interaction via commands and direct manipulation. In the near future, more seamless integration of Internet agents with other office applications will undoubtedly occur. The agents would have the capability to relieve users of major repetitive tasks. For example, every morning, based on the list of activities, the Internet agent could gather, organize, and summarize information for meetings. While a user is writing a report or a letter, the agent may automatically, without an explicit command from the user, search for relevant information.

Desirable Features in Internet Agents

A working balance in the following three factors is essential for the success and popularity of Internet agents:

- *Autonomy:* Autonomous agents will be able to make simple, and eventually complex decisions with minimum guidance from the user. Thus the agent would serve as the user's personal assistant without the hassle of explaining every detail.

- *Simplicity:* These agents need to be simple enough for even a novice Internet user to feel comfortable in using.

- *Economical:* Cost of the agents is an important issue as well. If the price gets too high, personal-computer users will be excluded from the market and only businesses would be able to afford Internet agents.

Some other desirable features for Internet agents are:

- *Believability.* Users tend to trust an agent to perform a task if the agent shows emotions. For example, joy or distress should be expressed when an agent's goal is successful or falls short.

- *Robustness.* Robustness is how well the agent stands up to the unexpected. The agents should be able to handle situations that, quite simply, have not happened before.

- *Standardization.* There must be enough standardization in Internet agents that different agents can interact with each other and be able to use the same information appliances.

- *Sociability.* Agents should communicate and share their resources and experience for similar tasks. Agents with appropriate social interaction will move the Web from the present client-server paradigm to a distributed and cooperative paradigm.

- *Learning and Cooperation.* An increased learning and cooperation ability is necessary for users to be able to train their Internet agents by using examples. In addition, the agents may learn from the experience of other agents on similar tasks.

References

1 Franklin, S., and A. Graesser, "Is It an Agent or Just a Program?: A Taxonomy for Autonomous Agents," *Intelligent Agents III: Proc. International Workshop Agent Theories, Architecture, and Languages,* pp. 21-35, Springer Verlag, 1997.

2 Petrie, C., "Agent-Based Engineering, the Web, and Intelligence," *IEEE Expert,* Dec. 1996.

3 Russel, J., and Peter Norvig, *Artificial Intelligence: A Modern Approach,* Englewood Cliffs, NJ, Prentice Hall, 1995.

4 Belgrave, M., "The Unified Agent Architecture: A White Paper," URL: http:\\www.ee.mcgill.ca\~belmarc\uaa_paper.html, 1995.

5 Hayes-Roth, B., "An Architecture for Adaptive Intelligent Systems," *Artificial Intelligence: Special Issue on Agents and Interactivity,* vol. 72, pp. 329-335, 1995.

6 Etzioni, O., "The World Wide Web: Quagmire or Gold Mine?" *Communications of the ACM,* pp. 65-68, November 1996.

7 Higgins, K., "Your Agent Is Calling," *Communications Week,* pp. 45-49, August 5, 1996.

8 Gold, S., "Autonomy Claims 10,000 Users of New Web Agent Technology," *Newsbytes,* August 13, 1996.

9 "Wireless: RadioMail Corporation Unveils Internet-On-The-Air," *EDGE,* pp. 17, June 10,1996.

10 "Internet Access: Tierra Highlights," EDGE, pp. 13, August 26, 1996.

11 Etzioni, O. and Daniel W., "A Softbot-Based Interface to the Internet," Communications of the ACM, pp. 72-76, July 1994.

12 LaMacchia, A., "Internet Fish," *AI Technical Report* 1579, July 1996.

13 Kiniry, J., and D. Zimmerman, "A Hands-On Look at Java Mobile Agents," *IEEE Internet Computing,* vol. 1, no. 4, pp. 23, July-August 1997.

14 Kay, A., "User Interface: A Personal View." *The Art of Human-Computer Interface Design*, edited by B. Laurel, Addison-Wesley, Reading, MA, 1990.

Works Consulted

Foner, L. "What's an Agent Anyway? A Sociological Case Study," http:\\foner.www.media.mit.edu\people\foner\Julia\Julia.html, 1994.

What You Need to Know About Internet Security as Internet-Related Commerce Drives Business into the 21st Century

Paul J. Plummer
President
Cytware Corporation

Nearly everyone knows that the explosive increase in the number of computer users and corresponding growth in the use of the Internet continues to dramatically change our world for the better. Electronic commerce over the Internet is the future and the driving force for potentially enormous business growth into the 21st century. Over $500 million in consumer purchases and $600 million in business-to-business commerce transactions were conducted on-line in 1996. By the year 2000, consumer purchases over the Internet are estimated to reach $7.7 billion.[1] The projected value of goods and services traded between companies is predicted to rise to $327 billion in 2002.[1]

Unfortunately, many computer users are beginning to learn the hard way about some of the potentially negative consequences of working and playing on the World Wide Web. Personal privacy and control of our computers have never been at greater risk. The enormous potential of Internet commerce cannot be realized until these complex security issues are addressed and resolved. The good news is that Internet and network users can protect themselves now by understanding the nature of these risks and taking precautions appropriate to their respective needs.

With Internet commerce between companies already headed for $8 billion in 1997, up 1,000 percent from 1996, Forrester Research, Inc. looked at which industries are at the center of the dramatic growth.[1] The three largest business types identified were manufacturers, services and utilities, and suppliers of office supplies and other computer-related supplies.[1] Manufacturers, such as makers of goods like small electronic and airplane parts, represented 38 percent of the total, amounting to $3 billion.[1] Service and utility companies also did about $3 billion worth of on-line trading, followed by suppliers, for a total of $2 billion so far this year.[1]

Although inter-company commerce over the Internet is dominated today by manufacturers, that will begin to change quickly as the efficiency of the Internet proves irresistible to companies that appreciate the ability to more fully describe their products, better configure orders, and offer round-the-clock (24x7) accessibility. Durable goods makers will lead the way, notching $99 billion in sales by 2002.[1] Non-durables will lag behind at $17 billion, most of it in paper, plastics, and apparel.[1]

It is well known that corporations, government, academia, and most major organizations have been using computers for years. What is news is the enormous growth in the use of computers by consumers. In the United States, 46 percent of American adults and teens have at least one home computer, and more than 31 million Americans have purchased a second or third computer.[2] More than 40 percent log 20 or more hours a week on their computers.[3] These connected households are also heavy Internet users. Women now account for 47 percent of Internet and on-line service users, compared with just 36 percent a year ago.[4] One of the recurring differences between U.S. and European users is age, with U.S. users averaging 33.9 years and their European counterparts averaging 28.8 years of age.[5]

Growth in the use of the Internet is explosive, conservatively estimated at the beginning of 1997 to be over 50 million people worldwide.[6] According to a recent survey by CommerceNet/Nielsen Media, 23 percent of all persons over 16 years of age in the United States and Canada have used the Internet during a recent month and still have access today. The report goes on to say that of all Web users, 15 percent (approximately 5.6 million people) have used the Internet to purchase a product or service on-line.[6]

The primary reasons cited by Americans for using their PCs are as follows: accessing the Internet (83 percent); working at home (78 percent); getting news and information (75 percent); keeping in touch using e-mail (59 percent), and performing school-related activities (59 percent).[7] While not currently one of the primary uses of personal computers, shopping may be the most obvious way the Internet is beginning to change our world. MasterCard International forecasts that Internet consumer purchase volume will jump from U.S. $800 million in 1996 to U.S. $10 billion in the year 2000.[8] Many of these users are buying everything from clothes and cars to homes. As travel agencies know only too well, many of these households are also making their own travel arrangements over the Internet.

The attractions for businesses conducting electronic commerce over the Internet are infinite, primarily focused on expanded commerce and more efficient operations. In generating billions of dollars through Internet trade, businesses are simultaneously cutting costs, reducing order processing time, and improving information flow. While expanding trade over the Internet, most firms are realizing a marked decrease in phone and fax use, allowing salespeople to concentrate on proactively managing customer accounts rather than serving as information givers and order takers.

What Are the Privacy Risks and What Causes Them?

Despite the tremendous opportunities for commerce, facilitated communications, and many other benefits offered by the Internet, many users are nervous—and rightfully so. The risks of playing and working over the Internet are very real. In the October Report to Congress by John Deutch, Director of the Central Intelligence Agency, Deutch stated that "In terms of overall threat to the U.S., the danger of computer based 'cyber' attacks is second only to that posed by nuclear arms and other weapons of mass destruction."

Currently, computer users connected to the Internet no longer have complete control over their own computers, which are open to intrusion ("hacking") without their knowledge or permission. Individual users and businesses are equally susceptible to unauthorized or illegal entries. Financial, personal, and business information can be read, altered, or destroyed by unscrupulous individuals. Until now, very little protection has existed for computer users in their homes and businesses.

Computer security breaches cost organizations millions of dollars each year. A recent survey of Fortune 1,000 firms produced striking evidence of serious problems in many commercial organizations.[9] Nearly half of the 205 firms that responded admitted that their computer networks had been successfully attacked and penetrated by "outsiders" in the past year.[9] Losses and associated costs were estimated at over $50,000 each by 84 percent of the respondents.[9] Moreover, 41 percent indicated losses of more than $500,000 per intrusion, with six of these companies estimating losses at over a million dollars.[9] Unfortunately, most of the more damaging unauthorized entries go unreported because most targeted organizations refuse to discuss incidents publicly or even report them to authorities for fear of publicizing the fact that their data is not secure.

The types of loss range from intellectual property theft, such as loss of trade secrets, to on-line bank robbery. The most notorious example was masterminded by a Russian hacker who used a laptop computer in 1994 to funnel $10 million in cash from Citibank branches worldwide to his accomplices' accounts in California, Israel, Finland, Germany, the Netherlands, and Switzerland.[10] Although most of the money was recovered and the Russian is awaiting trial in a federal prison in New York City, this case—the only publicly disclosed example of bank robbery in cyberspace—represents what crime is going to be like in the 21st century.[10]

Corporations are not alone in being attacked. When the Senate Subcommittee on Investigations held hearings recently on security in cyberspace, among the scariest findings was one involving the U.S. Department of Defense. There were more than 250,000 attacks on Department computers last year, and 65 percent were successful.[11] Little is known about who launched them, why, or what they found. Government agencies, academic institutions, and corporations are hit every day. In another widely reported incident, FBI investigators, armed with a court-ordered wire tap and a sophisticated computer program, tracked down an alleged hacker who had compromised computer networks at many sensitive sites including Harvard University, NASA, and the Los Alamos Naval Laboratory.[12]

PC users also have reasons to be alarmed. Although the intrusions into their homes are less dramatic, and much of the transfer of data is seemingly harmless, loss of privacy has become an issue of international debate. Some of this activity is completely innocent, such as sending product registration information to the software producer or informing a Web site owner about the people who are viewing or interacting with the site. For instance, many Web sites automatically know the user's e-mail address, what kind of browser is being used, what kind of computer is in operation, what pages in the site the user views, and how often and how long the viewer visits each site.

A common practice for the owners of these sites is to create a profile of the user's activities, which are then stored in a text file (known as a "cookie") on the user's computer. The next time the user visits that site, it will know how to direct information to the user. Most of this information is gathered invisibly, usually without the knowledge or consent of the consumer. With information sent to identifiable user addresses, the technology becomes the perfect vehicle for tracing and monitoring users. Now advertisers can keep track of exactly who is looking and target their advertising accordingly.

While much of this exchange is harmless, it usually occurs without the user's knowledge. America Online (AOL), one of the most popular Internet service providers (ISPs), collects information about its 8 million subscribers' incomes, lengths of residence, ages, and children. Then it aggressively sells its subscriber. AOL's on-line members are targeted increasingly by junk mailers, as are members of other on-line service providers, because AOL is not alone in this sales practice.[13] Member data from numerous organizations and commercial interests is extremely valuable and is currently being compiled, combined, analyzed, and sold with little or no legal restrictions. Many people insist that at the very least, these data exchanges constitute a serious invasion of consumer privacy.

Unfortunately, the same technology that allows these seemingly innocent, albeit possibly annoying interactions to take place is also being used for other, more sinister purposes. Thousands of toolkits for hackers are available on the Internet through over 1,000 Web sites. Furthermore, instructions for hacking are exchanged daily through e-mail. Those motivated hackers who are interested in malicious activity (known as "crackers") can hack their way into almost any PC.

The most well known example of malicious computer intrusions is the computer "virus." Viruses, which are bits of software code that overwrite or attach to programs and replicate themselves, are attached to e-mails, stored in documents on file servers or downloaded from the Internet. When the infected programs are executed, the embedded virus is executed too, thus propagating the "infection." This activity normally happens invisibly to the user. However, after propagating silently for awhile, a virus usually starts doing other more irritating things like writing messages on the terminal or playing strange tricks with the display functions. Despite increased awareness of the virus epidemic and increased use of different types of anti-virus programs, viruses remain a serious problem. The National Computer Security Association (NCSA) reported in the 1997 Virus Prevalence Survey that 33 out of every 1,000 machines were infected with a virus in any given month, compared to 10 out of 1,000 in 1996.[14]

While many consumers and commercial users know about viruses and take precautions through a variety of commonly available virus prevention software and network system deterrence protocols, viruses are unfortunately just the tip of the iceberg in terms of damaging activity.

The Next Generation of Computer Vulnerability

An example of an emerging, more sinister activity is commonly experienced by both commercial and personal PC users. This results in the loss of passwords and data, or in other instances, crashed hard drives. These problems are often caused by "Trojan Horses," which often appear as a free gift or free tool attached to e-mail, but have an alternative, malicious purpose. Viruses are often confused with Trojan Horses, which can start as though they are viruses but actually become Trojan Horses. Unlike viruses, Trojan Horses don't replicate themselves. Once downloaded, some of these intrusive program scripts snoop around for data and activate the user's Internet connection to send it back to the creator. They can send financial data from Quicken or other financial files, information about PC configuration, or nearly anything the creator wishes to retrieve. America Online users have recently been plagued by Trojan Horses, including more than three dozen that sent users' passwords to crackers.[14]

Officials at the Federal Department of Energy's virus hunting team reported another frightening example of this type of cracker activity. Earlier this year, they discovered a program called AOL4FREE.COM that arrives attached to an e-mail message. When activated, this Trojan horse quickly erases a computer's hard drive. Bill Orvis of the Energy Department's Computer Incident Advisory Capability team said "When it finishes deleting all your files it prints an obscene message on your screen."[15]

AOL and the National Computer Security Association (NCSA) recently warned subscribers about the danger of downloading e-mail attachments from strangers and AOL implemented a pop-up screen that warns users whenever they attempt to download an e-mail attachment that might contain a Trojan Horse. They also posted a software program that scans for thousands of Trojan Horses—but not all of them.[16] New Trojan Horses are being created faster than the software being developed to scan for them.

Other types of damaging cracker tools and techniques are "hostile Java applets," "logic bombs," "spoofing," and "worms." Java is a programming language in growing use on the Web. Applets are tiny software programs that can add features such as stock tickers to Web pages. A hostile Java applet can be distributed as an attachment in a World Wide Web document and executed either by Sun's Hotjava browser or Netscape Navigator Version 2.0 or later. An applet is not as dangerous because Navigator and other browsers restrict the functions of an applet, which runs inside the browser. However, a full hostile Java application runs outside the browser, so it has the most potential for doing damage. Once downloaded onto the user's computer, a hostile applet can damage computers in a variety of ways, from deleting files to crashing hard drives.

"Logic bombs" are codes that surreptitiously are inserted into an application or operating system and then cause it to perform some destructive or security-compromising activity whenever specified conditions are met. Logic bombs most often activate by predetermined dates. "Spoofing" is disguising the source of a program, a kind of impersonation of a valid user or computer. It is used to gain authentication where it would not ordinarily be given. For example, if crackers tell their own computers to use an IP number that is part of another corporation's network, then the crackers could have access to that corporation's computers. They may be able to read data, retrieve files, install viruses, or perform most other forms of malicious cracker activity.

"Worms," which are programs that propagate themselves over a network, reproduce themselves as they go. Consequently, these programs cause damage as they travel. One of the most notorious cases resulted in 1988 when a college graduate student at Cornell University unleashed a worm and caused the entire Internet to go down worldwide in 12 hours. The worm also crashed or severely crippled thousands of computers at universities and military sites.[17] Unfortunately, this potential vulnerability still exists.

Internet related scams to get financial information or entice computer users into sending money are often associated with other types of common e-mail invasions. Last spring, a spokesperson for Sprint said that e-mail messages allegedly from "Sprint Online" were sent to some customers requesting credit card and billing information. The fraudulent messages stated, "Problems with your account—please read—Sprint Online." Fraudulent attempts to gain personal information such as credit card information, Social Security numbers, and even telephone calling card numbers are on the rise, according to the National Fraud Information Center, which monitors reports of fraud on the Internet. The center said that con artists are also sending an increasing number of mass e-mailings in which they solicit money for illegal pyramid and other "get rich quick" schemes.[18]

Another common scam on the Internet causes consumers to receive expensive phone bills for calls they do not realize they have technically made. A good example of this scam occurred in January of 1997, when Web surfers received phone bills for thousands of dollars worth of calls they had not dialed. While they were downloading a "free" viewer program advertised on a Web site, international calls were being made connecting them to the "free" program, which was being operated in Moldova. The "free" program was a pornography site. AT&T inadvertently now has the names of all the individuals who viewed the sexually oriented Web site and the pornography program, how long they were on-line, and how often they went back to that site. Additionally, the Web surfers were reportedly being held responsible for paying the international telephone charges, given that the calls were in fact made.[19]

This example raises another problem that has everyone from parents' groups to the White House scrambling—the issue of protecting children from offensive Web sites and language, as well as unsupervised chat sites and similar inappropriate places. Although there are some software programs available to parents for filtering out many of these sites, it currently is extremely difficult for parents to guard their children from all of these risks.

The fragile nature of e-mail security is further evidenced by the increasingly common practice of forging e-mail, usually as a follow-up to the theft of passwords. The media reports of these typical occurrences include the defamation of character for victims, as well as lawsuits. Crackers or disgruntled employees write to the e-mail files of their victims and plant information on the victims' computers. The information is then attributed to the victims. Politicians, corporate managers and most individuals with any public visibility, including journalists, are learning the hard way that it is vital to have a secure e-mail protocol in place.

One example of forged e-mail is evidenced by a case earlier this year in which a former employee of Oracle was convicted of perjury and falsification of evidence used in an earlier $100,000 sexual harassment settlement against Oracle's CEO, Larry Ellison. The employee allegedly forged an e-mail from her direct boss to Ellison saying that she was fired 'per your request.' The e-mail was incriminating enough to force a settlement.[20]

Forging Usenet posts is also easy and potentially very damaging. Someone posted dozens of forged racist messages over Usenet to smear another user's reputation. Although the messages were canceled and the victim posted repudiations, the articles are still archived in several search engines and appear when the victim's name is a search target.[20]

What's Being Done to Protect Computer Users and How Successful Are These Efforts?

In the face of these threats, educators, legislators, government officials, corporate security providers, and a host of other groups worldwide are scrambling to protect us, with questionable success. Various strategies are being implemented to provide varying levels of PC, network, and Internet security. The explosive growth in the computer security industry is a testament to the seriousness of these new international problems. Unfortunately, despite the increased attempts to educate consumers, pass legislation, and install software applications and protocols, these efforts for the most part are complex, costly, and, all too often, ineffective.

Consumer education usually takes the form of warnings about the problems, although these efforts do not offer any real solutions. For example, in the wake of a controversy stirred by crackers who claimed they found a way to use the Active-X code to make bogus money transfers, Microsoft Corporation announced the formation of a program to educate users about the security risks of executable code. It should be noted that Microsoft was not able to correct the problem, only admit its existence.[21]

State and federal legislators and international governments and organizations have held numerous hearings, commissioned studies, and passed legislation and protocols in an attempt to protect users of the Internet. The problems posed by these attempted "solutions" are as varied as the security problems they address. When states have passed legislation, federal courts have nullified some of this legislation on the grounds that states cannot restrict commerce outside of their own jurisdictions. Debates abound over the Constitutional right to free speech and free trade, juxtaposed with the need to protect privacy and intellectual property.

Liability risks are also at an all time high, which is partially evidenced by the incredible growth in the number of lawyers now specializing in the area of computer security law. Consider the consequences to a doctor if sensitive medical information were broadcast illegally over the Internet. Multimillion dollar lawsuits might result. In short, there is no agreement on appropriate strategies for security, no way to enforce strategies, nor is there any territorial control over the Internet. The World Wide Web has no boundaries.

In the face of the serious threats posed, the computer security industry is booming. Computer security, little more than a university-based research issue a decade ago, is today an industry with revenues as high as $6 billion a year, projected to reach $13 billion by the year 2000.[22] New software products developed to address single types of security threats and privacy invasions are being developed and marketed internationally. Security firms now specialize in assessing, designing, procuring, and implementing a range of security solutions, including firewalls, combined with scanning and monitoring tools.

Firewalls are programs put into a system to make sure that the users cannot do any damage and crackers cannot access the network. The construction of a firewall typically is an inexpensive microprocessor-based UNIX machine with no critical data, with modems and public network ports on it and just one carefully watched connection back to the rest of the cluster. The objective is to protect a cluster of more loosely administered machines hidden behind the firewall. Unfortunately, firewalls by themselves do not protect networks and the computers in their systems from most malicious activity. According to the computer security branch of Information Technology (IT-SEC), the United Kingdom's government agency that certifies computer security products, the vast majority of commercial firewall protection software is faulty and unsafe.[23]

Cryptography, another important tool in computer security, refers to the practice of encrypting—or encoding—data so that only designated individuals can decode it. A system for encrypting and decrypting data is a cryptosystem, which usually employs algorithms for combining the original data ("plaintext") with one or more "keys." These keys are numbers or strings of characters known only to the sender and recipient. The resulting output is known as "ciphertext." The security of a cryptosystem usually depends on the secrecy of the keys rather than on the supposed secrecy of the algorithm. A strong cryptosystem has a large range of possible keys so that it is not as easy to just try all possible keys. Cryptosystems produce ciphertexts that appear random to standard statistical tests and resists many previous methods for breaking codes.

One of the problems with encryption is that the codes can still be broken. One example was announced on September 26, 1997 when Bruce Schneier, a cryptography consultant in Minnesota, announced that he had developed a Windows '95 screensaver that cracks encrypted e-mail messages.[24] His program works on multiple machines in parallel over a local area network (LAN). According to Schneier, his program can crack messages on a thousand machines in just 50 minutes. This program will only crack messages encrypted with RSA Data Security's S / MIME mail encryption standard and only messages that are encrypted with a 40-bit key. However, that is exactly the encryption that is being offered by the most commonly used versions of Netscape Messenger and Microsoft Outlook Express. Schneier also said that he released his program on his Web site to demonstrate the fundamental vulnerabilities in the S / MIME standard.[24]

Computer users dedicated to achieving secure transmissions are also using cryptography in combination with user identification devices known as "public keys," "private keys" and "smart cards." Many of these tools are used as part of an emerging electronic standard for electronic commerce known as the "secure electronic transaction (SET)" protocol. A number of companies are attempting to develop a variety of programs to run on SET standards, including point-of-sale applications, gateways, and a sort of virtual driver's license known as a digital certificate. These certificates allow computer users to verify the identity of the users on the other end of a transaction. Unfortunately, digital certificates have yet to gain widespread acceptance, so they are currently far from the goal of allowing electronic commerce to reach its massive potential.

Unfortunately, the average computer users seem to think that if they buy virus protection software products, they are safe. Clearly, computer security is a much more involved and complicated problem. With so many elements involved in computer security, even computer security experts are having a difficult time differentiating the products and appropriate strategies for securing computers. Even if each of these products and strategies were successful in delivering their promises, the complexity of sorting between dozens of different security software products and strategies is staggering, and usually very expensive.

The only real constant within this diverse computer security community is the broad-based goal of securing PCs and systems susceptible to everything from hackers and internal corporate crime to credit card fraud and other electronic commerce scams. Most of the products available to PC and network computer users merely protect computers from most viruses, which protect users from only a fraction of the potentially damaging activity resulting from the more significant cracker activity and fraud.

Although less well known than virus protection software, a number of security products are available to search the Web for hostile applets or scan inbound and outbound Internet e-mail, file transfers and Web-browsing traffic, such as Webshield. Another product, Groupshield, scans and cleans messages in real time within mail stores, private folders and public folders. These products tend to be limited to using time-sensitive, edited lists to do their jobs, allowing newer and unknown malicious scripts to pass. Other security efforts focus on taking the control out of consumers' hands by providing a service that automatically scans users' e-mail for known dangerous codes before the codes are delivered. Another software product, Webstalker, protects large corporate Web sites by monitoring all activities on a company's Web server. When a break-in occurs, the software kicks the intruder out of the system and notifies the company's system administrator.

Another product, Netcrypto, secures application-to-application communications by automatically encrypting all TCP / IP traffic between protected NT and UNIX servers. Other software applications and strategies attempt to make it possible for credit card payments to be collected on-line in a secure, automated transaction. One such product, Mallmanager, allows merchants to create a graphical interface for the collection of credit card numbers. That interface passes the customer information to Web servers, which then transmit it to credit card networks for authorization and payment.

Clearly, the complexity of sorting between the many different security software products and strategies and the multiple facets of the computer security challenges are staggering. Unfortunately, even computer security experts' strategies are often more like band-aids than real security solutions.

How Can Computer Users Protect Themselves?

Although PC and network users have varied needs, prudence demands that at the very least, every computer connected to a modem should be protected. This protection should include the following capabilities:

1. prevent unauthorized intrusion (hacks) from the Internet or other networks

2. prevent unauthorized reading of data on the hard drive

3. prevent unauthorized removal of data from the hard drive

4. prevent unauthorized writing to the hard drive

5. protect all personal and business information on the hard drive

6. prevent viral infections

7. hide files and documents

8. provide a written log of all uninvited intrusions and dispositions

9. provide application specific protection

10. allow the user to control a child's access to inappropriate Web sites and chat services

11. allow the user to set time frame parameters to control the amount of time spent on the Internet

12. provide a flexible and easy to use interface that empowers the user to control the flow of information

13. run in the background of the computer so there is no appreciable compromise to the computer's performance

Protecting the Personal Computer

Ideally, finding a single, all-purpose product that provides all of these protections dynamically is the best solution. One such product in the finishing stages of beta testing, Cytlok™, is designed to offer these features. This client-side program is designed to act as a sentry at the entrance to the computer, protecting the boot sector of the hard drive, allowing users to control the flow of information to and from their computers and protect user information from disclosure. Cytlok™ is also designed to monitor every program and alert the user whenever anyone attempts to compromise the integrity of the system.

Advancements like these in containing security breaches is client-side security, protecting the computer from not just viral infections, but also "hostile applets," "Trojan Horses," "worms," and "logic bombs." With such security measures, PC and network users should be able to regain the privacy and protection that is currently being eroded at an alarming rate.

1. Business Trade & Technology Strategies, "Sizing Intercompany Commerce," Forrester Research, Inc., Cambridge, Massachusetts, July 28, 1997.

2. Dell Computer Corporation, "Survey Shows Emergence of 'TechKnow Generation,'" September 9, 1997.

3. Dell Computer Corporation, "Survey Shows Emergence of 'TechKnow Generation.'"

4. IntelliQuest Information Group, Inc., "Latest IntelliQuest Survey Counts 51 Million American Adults On the Internet / Online Services in the Second quarter 1997," September 4, 1997.

5. "GVU's 5th WWW User Survey, Georgia Tech Research Corporation, 1996.

6. CommerceNet / Nielsen Media Research Survey, "Startling Increase in Internet Shopping," March 1997.

7. Holmes, "Trojan horses, hostile Java applets target home PCs," *USA Today*.

8. "Secure Electronic Commerce Takes a Major Step Forward in the U.S.," *Current Technology,* August 1997.

9. Warroom News Release, "1996 Information Systems Security Survey," November 23, 1996.

10. M.J. Zuckerman, "Security on trial in case of online Citibank heist," *USA Today*, September 19, 1997.

11. "A New Worry: Terrorism in Cyberspace," Intelligence Report, *Parade Magazine,* September 29, 1996.

12. Oppenheimer Wolff & Donnelly, "Law & Computers," Minneapolis, Minnesota, 1997.

13. David Braun, "AOL Snoops on Members' Privacy," *Privacy Times,* June 9, 1997.

14. Holmes, "Trojan horses, hostile Java applets target home PCs," *USA Today*.

15. "Trojan Horse," *Sacramento Bee,* April 18, 1997.

16. "Online Password Theft, Trojan Horse Warning," ABC 7 Online: WJLA-TV, Washington, D.C., 1997.

17. "Black Widows and Evil Worms," The Computer Network, May 5, 1996.

18. "Sprint Warns Customers About Internet Scam," *Kansas City Star,* April 27, 1997.

19. Simson Garfinkel and Gene Spafford, *Web Security and Commerce,* Ingram Book Company, May 1997.

20. "This PC Week," *PC Week,* March 5, 1997.

21. "This PC Week," *PC Week,* February 20, 1997.

22. M.J. Zuckerman, "Security on trial in case of online Citibank heist," *USA Today*.

23. Dan Jellinek, "Net security fails fire drill," *Internet Intelligence Bulletin,* August 20, 1997.

24. Simson Garfinkel, "S / MIME Cracked by a Screensaver," *Wired,* September 26, 1997.

Intranets and VPNs: Strategic Approach

Jai R. Rao
Adjunct Professor
Washington University

Abstract

The Internet has been accepted as the company network of choice based on its flexibility as an international standard. Companies face the challenge of how to use the Internet cost-effectively to support network security requirements and interoperate with legacy systems. Virtual private networks (VPNs) create secure wide-area network links over the Internet, making it possible for a company to connect its wide-spread manufacturing facilities. This paper addresses some of the strategic issues regarding the deployment, upgrade, and maintenance of intranets and VPNs.

Company Information Sources

The information–generating sources of most companies can be categorized as financial, management, human resources, plant facilities, and products. Each company's management information system (MIS) is composed of these essential information sources. Data regarding manufactured products is a small part of what is stored in the MIS; most product data is stored in hard copy. Recently, however, companies have begun storing their product data in soft copy, which requires the upgrade and expansion of their MISs. The MIS also stores data regarding services a company offers that are related to its products; for instance, an insurance company might need to store information concerning allied healthcare products. The emergence of the World Wide Web and the Internet has given rise to a multitude of means for information display and exchange formats. In the past, corporate repositories or data banks were static until a manual update could be applied because companies lacked the appropriate technology to make the updates automatic. Because companies fear compromising the security of their proprietary information, they have been hesitant to exchange the sensitive internal information over public networks. Strangely enough, not so long ago companies believed their sensitive data to be more secure on a magnetic tape in a well-packed container than on some telecommunications medium. Corporate information processing traditionally took place at the MIS center, and most product-related information stayed where the product was manufactured. Telecommunications technologies were used merely to exchange personnel information via proprietary e-mail systems with limited capabilities.

Today, product information is a business driver. Because of ubiquitous electronic distribution facilities, the geographic location of product information sources is virtually irrelevant. Similarly, accessing information depositories is no longer an obstacle to doing business. Corporate information sources are akin to the melting ice caps that flow down a mountain in the form of streams, which develop into rivers, lakes, and eventually merge into the oceans. These oceans of data become corporate, country, and world data banks. The communications networks are the foundation for corporate data repositories' data distribution and capabilities.

Every level of enterprise operation is powered by information sources. The majority of these resources are owned and managed not by the government, but by private companies. In general, information sources, whether private or public, are product sensitive, and their distribution must be controlled through appropriate security measures. If telecommunications technology is the primary carrier of this information, proper access control, authorization, and authentication features must be built into the infrastructure. As this technology continues to evolve, various cost-tradeoff options will be available for corporate networks; these options are explored in the following section.

Intranets and VPNs

For a long time, enterprise networks operated with a low–speed international standard CCITT X.25 and proprietary systems network architecture (SNA), DECnet-based wide-area networks for national and international information distribution. They were limited to Modem speeds with a maximum of 56 kbps. The high-speed telecommunications revolution started at the local area network (LAN) level with 10 Mbps, 100 Mbps, 1 Gbps and even higher speeds. Telecommunication speeds accessible to the user desktop, however, are inhibited by the workstation's internal bus speeds. It has been predicted that bus speeds will reach a speed limitation, and future generations of desktops will use crossbar switching to support workstation peripherals, which will replace the back-end bus.

In corporate networks, the current trend is to upgrade X.25 and proprietary packet-switching backbones to frame relay and / or asynchronous transfer mode (ATM) cell switching-based networks. Public carriers such as AT&T, MCI, and

FIGURE 1

A Typical Private Network

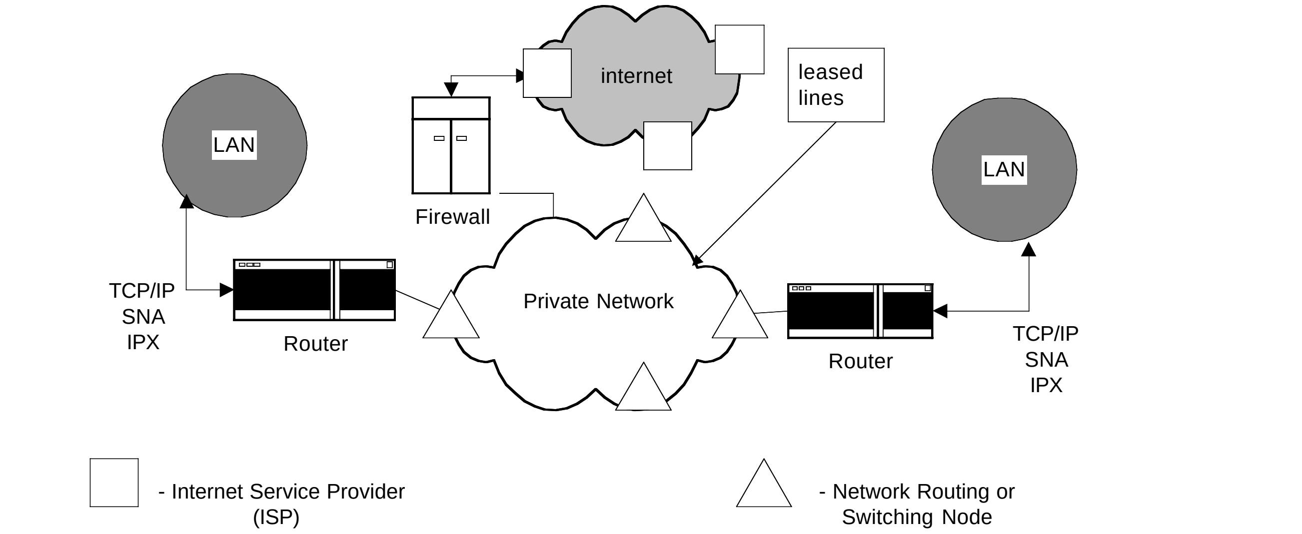

Sprint are offering new technology-based network services. Because of dedicated leased circuits for private networks, corporate network owners have more faith in the security of their information distribution systems. These private networks have primarily used connection-oriented network services (CONS), i.e., a virtual circuit setup that is dedicated to the user session for actual information exchange. These private networks have assured corporate management of the security of their proprietary information.

Figure 1 illustrates a typical private network. A private network may consist of a combination of leased lines, X.25 network, SNA network, and / or frame relay, ATM networks with assigned virtual circuits. However, the recent emergence of Internet protocol (IP) packet-switched networks, in which information packets travel on top of connectionless network services (CLNS), has created immense network security concerns. A firewall and a router can be used to filter incoming and outgoing packets to enforce access control specifications by the private network. In this network scenario, the private network is isolated from the public Internet via a firewall system. The public Internet can facilitate access to and browsing of the public and private databases via product Web pages. However, especially as relates to private networks, corporations must bear the burden of ongoing network-related issues:

(1) monitoring and resolving network problems
(2) upgrading the equipment for capacity and performance
(3) setting up redundancy
(4) provisioning new circuits
(5) enhancing software
(6) maintaining routing and switching tables
(7) staffing and training.

To relieve themselves of some of these responsibilities, companies have inclined toward public Internet services.

The networking industry has started zeroing in on network security issues, the Internet's weakest point, with special attention to intranets and VPNs. The public Internet provides a transport mechanism for interconnectivity across national and international boundaries. As it stands today, the public Internet provides logical functions at OSI Level 3 routing, OSI Level 3 switching, OSI Level 2 switching, and OSI Level 2/3 frame-relay capabilities. The technology trend has been to use a combination of OSI Level 2 and 3 functionality to gain performance and processing speed. Internet standards are leading the way for national and international network interoperability.

Intranets are another innovative approach for the use of IP. Simply stated, intranets allow enterprises to develop internal information sources and use search engines to expedite the access of corporate data. According to a Yankee Watch Data Communications white paper, the ten executives polled gave as many different definitions for "intranet."[1] Companies often started their intranet systems to enable the access of corporate-wide human resources information. This development has led to numerous other databases. As a result, users sometimes must spend considerable time finding the information they need in the maze of data warehouses. Poor planning and approach in designing intranets has created a waste of resources and non-productive searching cycles. An illustration of a typical intranet is shown in *Figure 2.* As shown, the public Internet is separated from the private IP-based network via a firewall system. The primary difference between *Figures 1* and *2* is that the private wide area network (WAN) based on CCITT X.25 and / or proprietary protocols is replaced by the standard IP. Similar to the public Internet, a private IP network has the capabilities of installing its own Web pages and browsing the capabilities of internal and external resources. Some examples include project schedules, budgets, employee address book, human resources, collaborative projects, workflow, and external market data, etc. Those who have been in this industry for the past 20 years remember what this capability means to productivity. Years ago, for example, in a commercial aircraft company, ordering parts from remote locations around the world often produced the wrong part, due to human error, and resulted in a long waiting period, which always frustrated the customer. With intranets

FIGURE 2

An Illustration of an Intranet

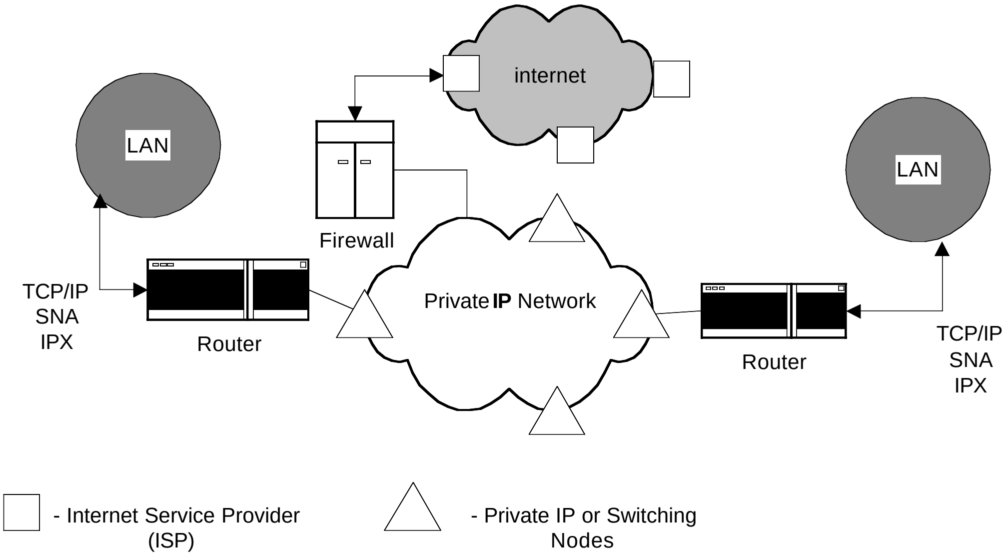

these problems can be easily eliminated or minimized considerably. Furthermore, with multimedia graphics capabilities, the chances of ordering a wrong part almost completely eliminated. In addition, intranet use facilitates the establishing of an internal corporate standard. This has helped enterprise-wide information sources exchange data at will. In essence, well-planned intranets can improve productivity and reduce duplicate databases, which lowers costs. Intranets have shown the cost benefits of using an industry standard protocol on their WAN. The immediate benefits of intranets can be achieved even with a simple implementation, while the IETF standards committees are addressing the long-term interoperability issues for Internet and intranets. The network security has become a major concern in the public Internet environment. Using intranets reduces corporate network overhead and improves productivity through enterprise resource sharing. Databases can be co-located in one place, and users can access these databases at will, irrespective of their geographic location. However, the issues regarding the maintenance a private network and establishing international points of presence are still a major cost concern.

Intranets and Internets have facilitated considerable improvements in network management and network functionality and may even help in building scaleable architectures. Maintaining a public Internet and a private intranet will contribute to the overall enterprise network cost. The private intranets are not ubiquitous, just like telephone service around the world. The industry is looking for a ubiquitous IP-based WAN service, with an underlying concern for information access and transfer security. This concern has lead to the implementation of enhanced networking capability, known as VPNs. The VPN concept is illustrated in *Figure 3*. As shown, in addition to the IP transport protocol, the packets are encapsulated in a point-to-point tunneling protocol (PPTP) with encryption and authentication requirements superimposed. Among the number of solutions proposed, one evolving standard known as level 2 forwarding (L2F) offers considerable promise. The L2F is expected to provide a secure, encrypted tunnel between network nodes.

According to several industry sources, IBM employed a router-based approach in deploying its VPN capabilities.[2,3] IBM routers will consist of TN3270e VPN software. The SNA-based networks can create a VPN that encapsulates SNA traffic without requiring costly leased lines. The tests have shown that using encapsulation is much cheaper than converting SNA to IP and then back to SNA. Intel is also incorporating VPN capabilities in its Express Router. The Intel Express Router uses 144-bit encryption algorithm and is incorporated at data link level, where it will not require the use of PPTP. Others have taken the approach of installing a separate VPN box within a firewall using 40-bit encryption algorithm and Microsoft point-to-point compression (MPPC). In all these approaches one key design consideration is: how many simultaneous connections are supported?

If a company requires it, its total IP network spans can be partially or completely encrypted and authenticated per the VPN specifications. The VPN makes it a secure network while still providing all Internet access facilities. One needs to be concerned about performance and cost trade-off in deploying VPN ubiquitously; many businesses may need ubiquity of Internet service, and may not require total VPN capabilities. The VPN based on the Internet may still offer certain cost savings through (1) outsourcing the responsibility of backbone network capacity planning and upgrading, (2) no immediate point of presence infrastructure cost, (3) scalable network capacity and cost, and (4) shared network maintenance and support costs. VPN capability is the key reason enterprises are flocking to the Internet despite the concerns about performance and instability.

Company Strategies

Internet vendors are offering a variety of deployment options for intranets and VPNs. Some of these options include upgrading the existing equipment or replacing an entire system. For instance, some vendors offer the following capabilities for the router-based networks:

FIGURE 3

A Typical Virtual Private Network (VPN)

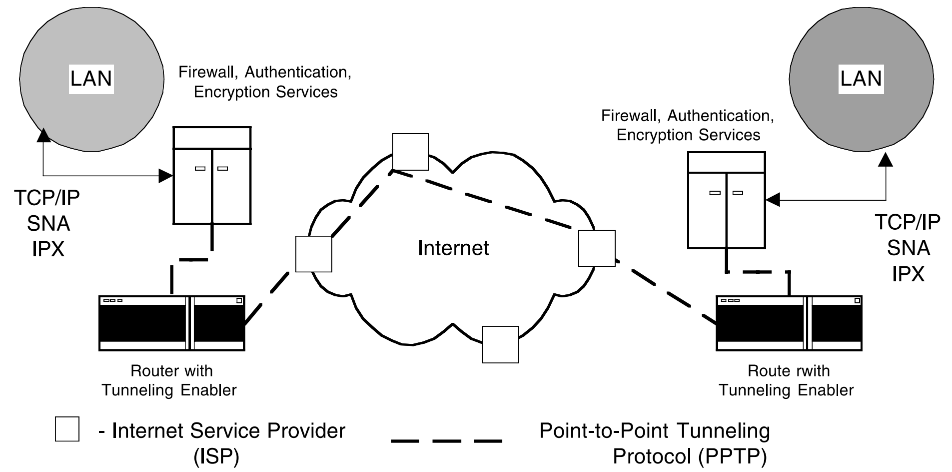

- network address translation, end-to-end security via a firewall, route authentication, enveloping, and encryption

- IPX-to-IP gateway for NetWare (also called IntraWare) LANs to interoperate with Internets

Other solutions are based on using a dedicated processor situated between the router and the channel service unit / digital service unit (CSU / DSU) to provide encryption, compression, and authentication. Some examples of intranet implementations include the Web browser access to mainframe applications, as an extension of the legacy client-server functionality. This capability has also minimized the controversy over distributed versus centralized processing. The challenge now is discovering how a browser can access mainframe applications. Today, many gateway products are being offered for Web access to the host. The first step would be to convert the legacy terminal data stream into a browser compatible language such as HTML or Java. *Figures 4* and *5* provide a simple illustration. Here the Web server consists of conversion software. The key design consideration is: how many simultaneous conversations can a gateway server support? If the gateway is locally attached to the host then intranet / VPN needs to provide network security to satisfy the client needs. Other pertinent considerations in setting up access from a Web to a mainframe are (1) Web server platform, (2) end-user data and application access requirements from the Web, (3) security and encryption requirements from the gateway to the host and the browser to the gateway, and (4) simultaneous sessions to the host.

As a starting point for building intranets, the following steps are suggested:

- identify the communities that exchange information
- identify what information is being exchanged

FIGURE 4

Web to Mainframe Access Gateway

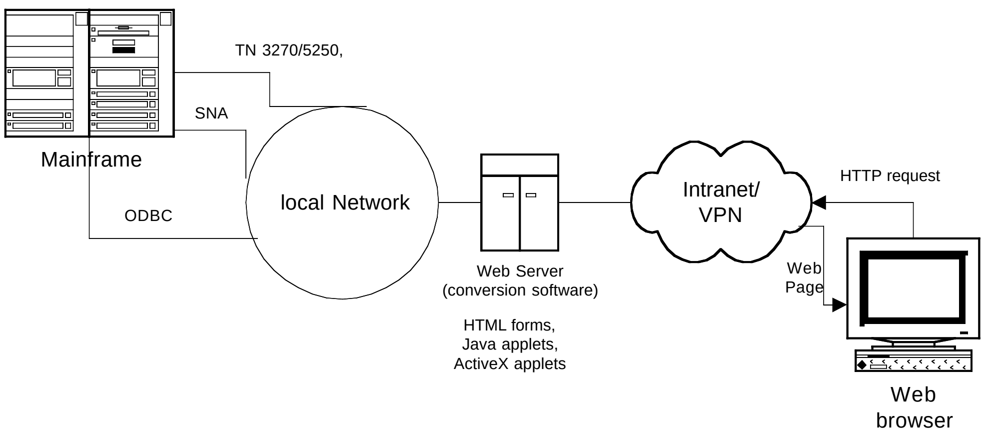

ODBC - Open Data Base Connection
HTTP - HyperText Transport Protocol
VPN - Virtual Private Network

FIGURE 5

Intranet in the Web-to-Host Access Application

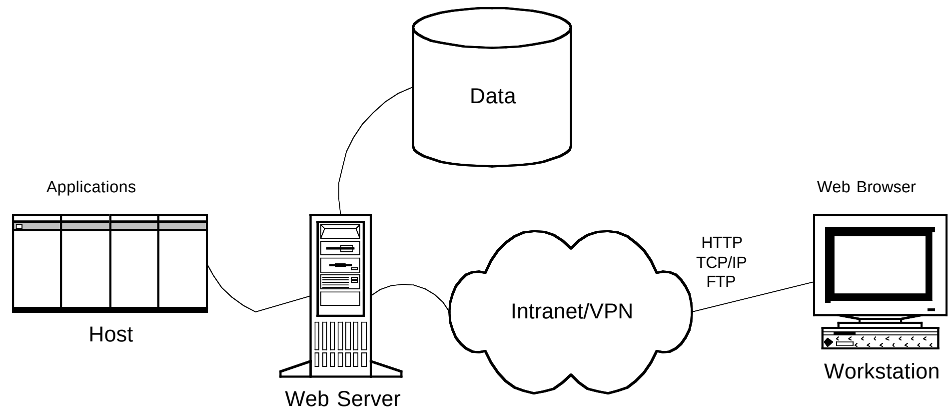

- identify what types of information are ready for on–line access and prioritize them
- identify the current systems and the cost of creating new ones
- re-prioritize your list based on cost and ease of implementation.
- plan overall architecture
- develop applications, create and convert documents, sell the intranet capabilities and the whole process in-house, and train users

Looking at today's enterprise environments, applications of the Web to mainframe access is still a requirement of 70–80 percent. The enterprise-related information sources—customers, manufactured goods, services, sales, etc.—are distributed all across the world, as shown in *Figure 6*, and the network infrastructure is depicted in *Figure 7*. Corporate repositories—policies / procedures, investor relations, e-mail, and announcements, etc.—are kept up to date at the headquarters level. The volume of interaction with headquarters is considerable. *Figure 7* illustrates a real world example of a private network with a frame relay–based public backbone network. This scenario might be true for a corporation challenged with adapting to intranet and VPN to meet the requirements of providing World Wide Web capabilities to its customers and employees.

Each of the geographic locations identified in *Figure 7* is a manufacturing facility for a specific product. The operations cycle illustrated in *Figure 6* is to be superimposed onto the network infrastructure deployed in *Figure 7*. There will be product-specific home pages for each of these locations that will be available to the employees and customers with the rules of access restrictions applied. Similarly, there will be an enterprise home page with pointers to common universal resource locators (URLs) and to product specific URLs as shown in *Figure 8*.

FIGURE 6

A Typical Enterprise Electronic Commerce

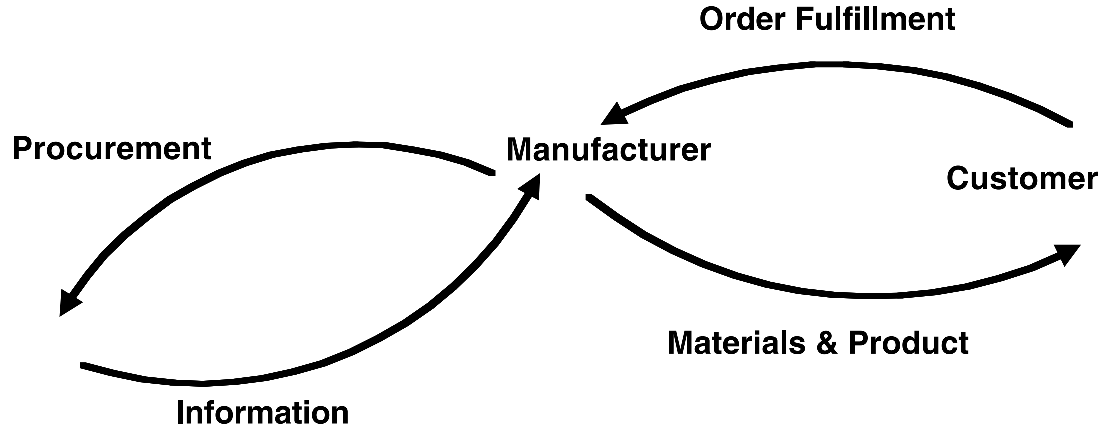

FIGURE 7

A Typical Private Network

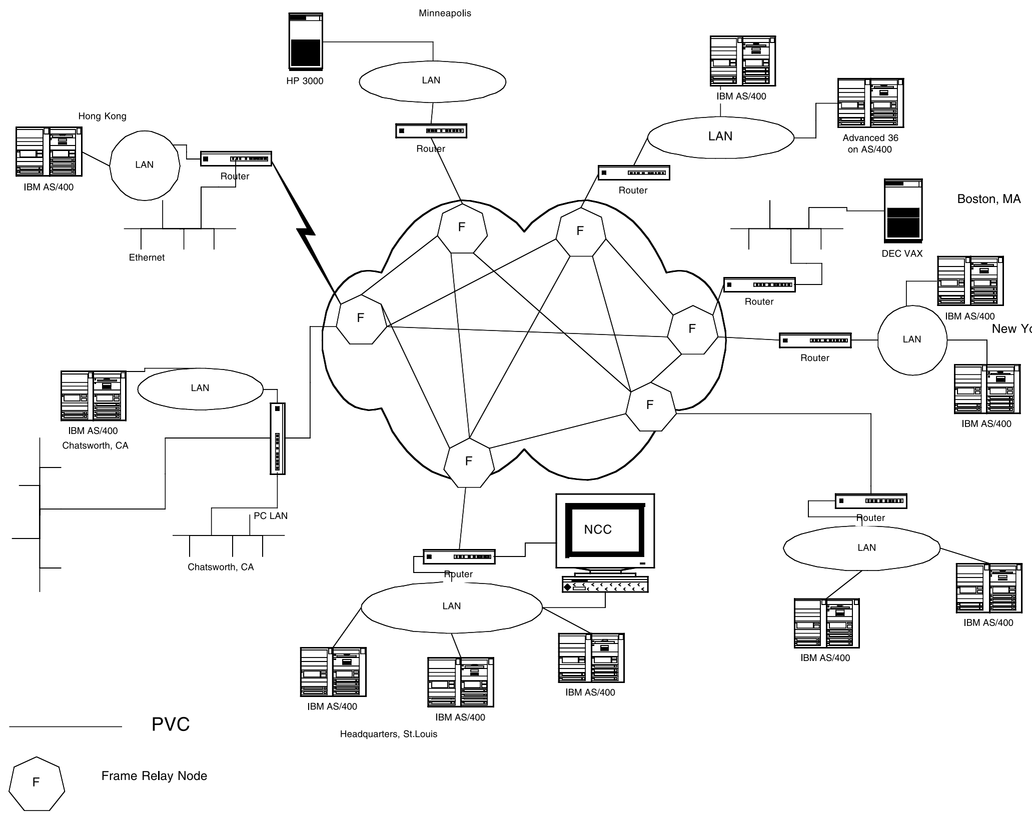

As has been discussed, the process of deploying intranets / VPNs for an enterprise incorporating its legacy systems is a lengthy one. It requires reengineering, performed by skilled professionals, to make it cost-effective.

References

1 *Yankee Watch Data Communications*, vol. 11, no. 13, December 1996.
2 *InfoWorld*, August 11, 1997.
3 *LANTIMES*, September 15, 1997.

FIGURE 8

Logical View of Corporate Resource Locator

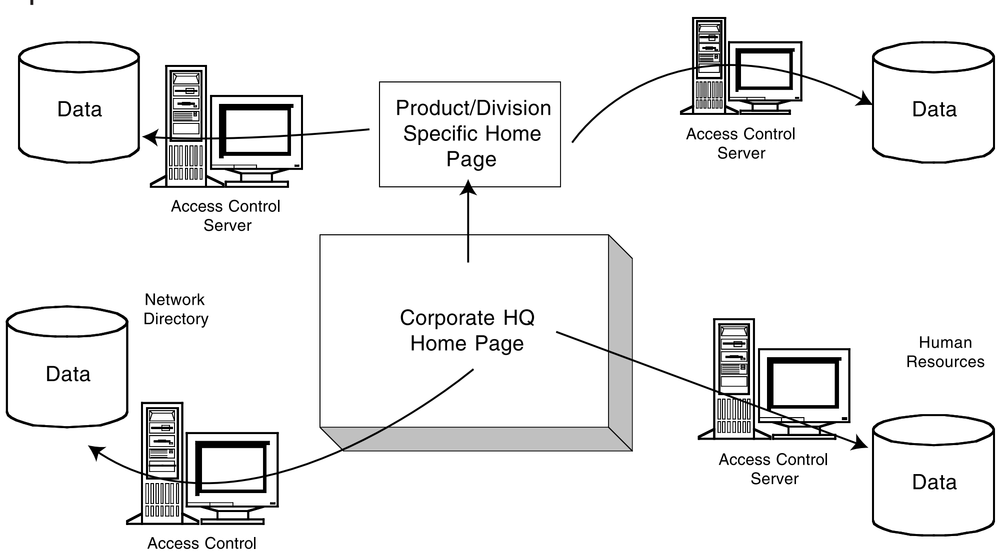

Integrated Voice, Video, and Data Services Over the Final 100 Meters of Connectivity

Michael J. Sodergren
President
Streamline Communications Corporation

Mike Heilbron
Applications Engineer
National Semiconductor Corporation

Abstract

The ability to support networked multimedia applications is becoming a mission-critical requirement for today's distributed, virtual enterprise. The push for networked video is the latest media form to impact network decision making in enterprises of all sizes. People are now interfacing globally through virtually all media. As such, the people's use of the enterprise's communications resources for data applications (e-mail, Web browsing, file and print sharing, etc.), voice applications (calls, voice mail, ACD, IVR, audio-to-text, etc.), and now video applications (broadcast, playback, real-time interactive, etc.) makes networking-investment decisions about the final 100 meters not just local, but global in nature. The customer's decision about video over the local area network (LAN) exemplifies just how much of a wide area network (WAN) issue a LAN decision can be. In fact, the various video over the LAN approaches addressed in this paper which are being refined by ITU Study Group 15 have a constant upon which they depend, the use of the synchronous/isochronous public switched circuit network (SCN) (a.k.a., the switched digital service, or SDS WAN).

As video is just now beginning to join voice and data in the realm of mission-critical communications assets of the enterprise, users are looking beyond just adding stand-alone video applications to their LANs. They are seeking the full integration and even synchronization of video with voice and data applications. Most analysts agree that this integration of services will traverse the LAN cable plant and infrastructure. Vital to the successful integration of video into the mission-critical communications-applications suite is the ability to match or exceed current end-user expectations for quality of service (QoS) for all of the previously stand-alone application services of voice and data. It is important to note that these expectations are based on years of user experience, not vendor/product capabilities.

A standards-based, vendor-independent model for high-quality/low-priced multimedia teleconferencing over the LAN of this new application medium is vital for mass success Solutions are required which leverage existing standards and which do not require expensive upgrades to existing hardware or software, an undepreciated infrastructure valued in the billions of dollars. This paper discusses the key issues to consider in determining enterprise-network solutions and will emphasize one standards-based network choice for the final 100 meters which delivers non-real-time and real-time interactive integrated services to the global enterprise multimedia desktop today without requiring end-user compromise.

Multimedia

Multimedia has invaded the enterprise-network marketplace, and users are demanding the integration of computer and telephony services. While users once were content to have access to stand-alone PC applications on client-server data-file-sharing networks and phone systems, this is changing rapidly. Users are now beginning to transfer and exchange many new types of content and even titles across their networks. Computers and telephony are being integrated to create a linkage between the parties of a "call" and with stored information which is relevant to the conversation. Also, during the past five years, due to increased emphasis on productivity in the workplace and, social and political turmoil in different parts of the world, videoconferencing has become established as a new means of communication for users in the enterprise environment. As a result, users have become accustomed to the capabilities and limitations of each of these independent applications and services. They are using the Internet, voice mail, interactive voice response (IVR), automatic call distribution (ACD), shared whiteboard/ applications, and videoconferencing-enabled applications more and more. At the same time, the world outside the enterprise continues to evolve into a new model. Companies are also becoming more specialized and distributed or "virtual" in nature.

What users in the networked enterprise are ultimately seeking is the ability to have the continuous and synchronized presence of all media forms (voice, video, audio, data) at the desktop. The PC should be able to support a videoconference in co-existence with broadcast TV, audio/video playback, file transfers, and real-time data-collaborative computing. Users should be able to make their windows to the world look like the PC screen in Figure 1 or any combination of their choosing.

FIGURE 1

Presence of All Media Forms on the PC

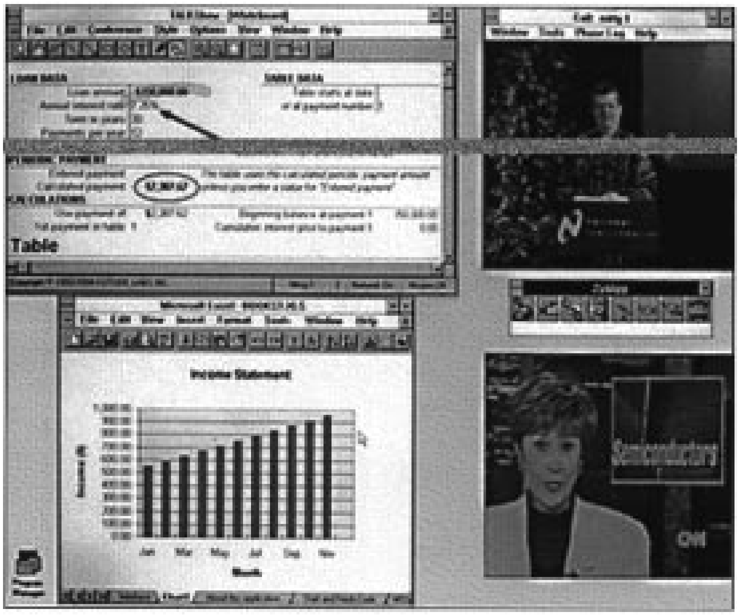

The Distributed Global Enterprise Need for "Connectedness"

One of the most fundamental changes over the past few years is in the very nature of how the enterprise itself interfaces with the outside world. Previously, the boundary of the enterprise was fairly clear. Companies had internal processes, procedures, and information and services infrastructures which were unique. As such, the internal networks of a company were also inherently unique. The boundary between a company's phone systems, LANs, videoconference systems, and the outside network (WAN) was fairly clear. Now, as organizations have evolved to a model where partners, vendors, customers, and other outsiders require access to the enterprises' phone, computer, and conferencing resources, the boundary between enterprises is becoming much less clear. For example, videoconferencing OEMs such as PictureTel and Compression Labs Inc. (CLI) found that it was a limitation to users and to sales growth to offer conferencing systems based on proprietary WAN communications-compression algorithms and call-signaling procedures. Users wanted to talk to anyone, anywhere, anytime, without regard for what brand of videoconferencing system they or others they interfaced with had purchased. As a result, the H.320 videoconferencing standard was born.

Integrated Voice, Video, and Data Systems

Today, phone systems, networked desktop computers, and videoconferencing systems are largely accessed by end users as independent "stand-alone" assets. *Figure 2* illustrates the distinct infrastructures of these three media types that is typical for today's communications systems. One reason why these services have not been integrated is that there are inherent physical differences between the networks for each of these services. From a physical network-infrastructure standpoint, each service is very different in its implementation. For example, in the enterprise environment, phones are typically connected to centrally located private branch exchanges (PBXs) or smaller key telephone systems (KTSs). The wiring for these systems tends to be cheap (Category 3 voice-grade cable) and may run for a kilometer or more between the desktop and the switches' line card. By contrast, LAN wiring systems based on the IEEE 802 specifications may use all sorts of cable and are typically limited to run lengths of 100 meters from the desktop to the HUB located in the wiring closet for structured wiring systems. Videoconferencing systems typically utilize yet another physical infrastructure by employing direct or inverse multiplexed WAN interfaces such as T-1s and PRI lines provided by local exchange carriers (LECs, RBOCs, PTTs, etc.) or interexchange carriers (IECs, CAPs, etc.).

While there are numerous other physical differences between the fabric that interfaces these three different networked resources, the total of these differences do not represent the biggest challenge in bringing about integrated services at the enterprise user's desktop. Research on pioneering users of networked video demonstrates that the addition of the new medium of video to existing networked systems is causing serious problems and a compromise to the quality of mission-critical applications. When given only a choice between being "satisfied" or "not satisfied" in a satisfaction survey on an assortment of multimedia applications, the Yankee Group reports that desktop videoconferencing was at the absolute bottom of the list of services users reported as being a "satisfactory" experience. Barely 10 percent were pleased with the current cross-section of offerings they had used. This result compared to a satisfactory rating of 49 percent for e-mail / fax and 66 percent for room-based videoconferencing.[1] *Figure 2* illustrates an example of a typical enterprise infrastructure for voice, video, and data services at a medium-sized enterprise location.

FIGURE 2

Today's Typical Enterprise Infrastructure

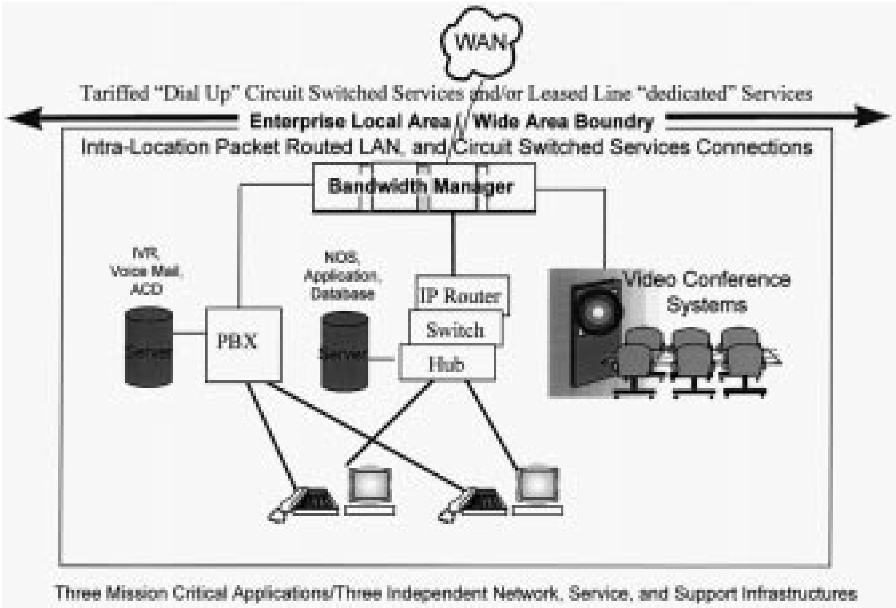

Two Mission-Critical Desktop Assets with a Third Yet to Come

Today, the end user in the typical corporate-enterprise setting has two mission-critical desktop communications assets: the networked PC, and the telephone. While the range of applications supported by these platforms varies in sophistication and the unique needs of the individual, norms have been established in all of the vertical (industries) and horizontal (job functions / professions), which have become, in effect, the competitive benchmark. For example, a marketing manager in the semiconductor industry could not live without a desktop phone system with voice mail and a desktop PC supporting e-mail. High-powered individuals in this profession would not even consider this to be the bare minimum for the job. This is but one example of the two mission-critical assets employees must have as baseline communications resources. Given the deployment of networked PCS in excess of 100,000,000 units and near ubiquity for phone sets with basic voice-processing applications such as voice mail, it is safe to lump all industry and job segments together and claim that the networked PC and the modern telephone system phone set are the norm as mission-critical communications assets for the enterprise knowledge worker. However, on a less ubiquitous scale, an increasing percentage of these workers are also adding non-real-time video and real-time interactive video as a third medium in the mission-critical mix of communications assets.

Today, however, the vast majority of video communications involve the worker physically leaving the personal communications resources of their individual work areas and going to a shared resource down the hall, out of the building, or beyond the premises. A classic example of this phenomenon is the room videoconferencing system. For this kind of communication, the employee (most often more than one) will schedule a time to use the shared resource with a conferencing administrator who not only schedules the conference based on availability, but also pre-engineers and may even set up the calls between the parties of the intended conversation(s). This is in marked contrast to the spontaneity with which these same people will pick up the phone to make a "call" or "fire off" an e-mail. *Figure 3* illustrates the nature of these mission-critical communications assets from the perspective of the end user at the enterprise desktop.

Interactive Multimedia Communications

Interactive multimedia proposes to bring together the best of the computer and telephony infrastructures to provide a new means of communication incorporating voice, video, and data that mimics real-time face-to-face interactions. What a concept—if only it were possible.

However, it can be. Do you want to place a videoconference call to an H.320 room system from your ISDN-connected desktop? Do you want to whiteboard or link applications over the Net? Do you want to set up multiparty phone calls to anywhere in the world on demand from your PBX-connected phone? Do you want real-time quality of service (QoS) and low cost of service (CoS) for integrated voice, video, and data services? There are many market claims about systems which offer the right CoS and QoS, but whose standards are we using to judge just how good, "good" is? I don't know what you expect during a videoconference call with an important client, but I have a distinct feeling that having a vapor trail of artifacts swirling about my head, accented by a dubbed-movielip synch has a overall negative impact on my content, credibility, and image. I know I want a better-quality video experience than this, but even that is not enough. What I really want are integrated-services solutions where voice, video, and data are seamlessly blended and synchronized to help me get my work done. This means, for example, that when important clients try to reach me via impromptu video calls to my desktop, I need my trusty voice-mail system to greet them after three rings if I am not there. This also means

FIGURE 3

Two Mission Critical Desktop Communications Assets with a Third Infrastructure for Video Conferencing (Currently NOT an Integrtaed Service Desktop Platform

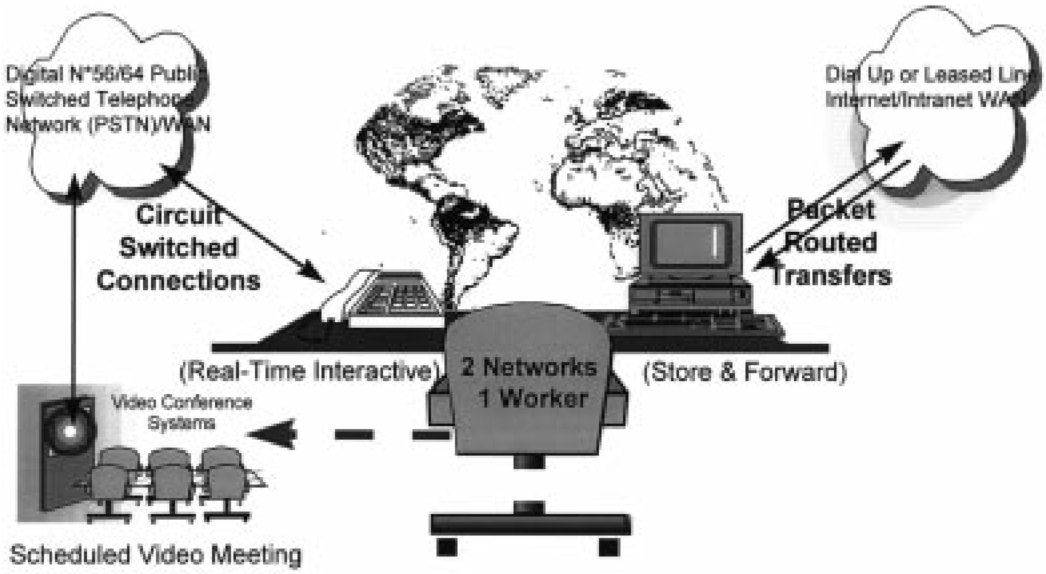

that integrated services should not result in any of my services receiving a priority downgrade. Important? Yes, and available today.

The Network Infrastructure to the Desktop—a.k.a "The Final 100 Meters"

To better appreciate the challenge to achieving integrated "multimedia" communications services at the enterprise desktop, one needs to look at how things are connected today. First, the network infrastructure for telephony, data, and video communications are separate assets optimized for their respective applications. While not universally the case, the vast majority of desktop data communications are encapsulated in packets which traverse cable plants connected to hubs and switches which are, by IEEE convention, located within 100 meters of the enterprise desktop. The physical-cabling plant is typically comprised of Category 3 or 5, four-pair, unshielded twisted pair (UTP) copper wiring. Internet protocol (IP) or IPX addressing is used to "route" information from source to destination on a "least-cost, best-effort" basis. The equipment of the data network within the enterprise (a.k.a., the LAN) is a distributed architecture, meaning that hubs, switches, routers, servers, etc. do not need to be centrally located and typically are not. By contrast, the phone system (PBXs, telephony processors, servers, etc.) tend to be centrally located within a premises or campus setting. The distance between the phone and any individual user's PBX line card could be thousands of meters, and the cable plant is typically very inexpensive UTP cable that is less than or equal to the LAN-cable grade.

Both voice and data communications systems are designed to allow enterprises to support locally secure and "free" communications while managing the interface to the public WAN. Videoconferencing systems have their own separate network, sometimes with a managed WAN interface and sometimes without. In rare instances (less than 50,000 systems deployed worldwide to date), additional networks specifically for carrying video traffic are provided to individuals at the desktop. This is in stark contrast to voice- and data-network services which are widely available at the desktop. In this case, a third cable plant may be required to the desktop. Use of multi-pair bundles may be an option for some administrators with well-planned, structured wiring systems, but the net effect is still three independent cable plants and systems that just happen to share the same sheathing for maybe 100 meters. Forgetting our lofty goal of integrating video with voice and data services, just the task of integrating voice with data services is seemingly a daunting task. Ultimately, the lowest common denominator in the equation is the physical layer or cable plant. By virtue of the needs of high-speed data networks, the distance between the desktop and network equipment is 100 meters for LANs. Therefore, any attempt to run distributed data communications over the PBX cable plant would simply not work for today's high-speed data communications. It is for this and other reasons that industry analysts such as the Gartner Group predict that the integration of voice, video, and data services will ultimately traverse the LAN cable plant.

Another key factor in network infrastructure relates to the nature of the communications that are traversing the infrastructure. In the case of the telephone, the communications are real time and interactive. As such, "circuits" are created for the duration of all "calls" which support full-duplex communications with minimal latency. In the case of data, the communications are typically not real-time interactive. Instead they are "store and forward" in nature. As such, routes are established and packets of information are transferred based on shared network availability. More than 99 percent of all LAN equipment in the installed base today is designed to support half-duplex communications. The point is that neither the local voice network or the local data network were designed to optimally carry the other's media form. If you are becoming discouraged at the prospect for integrated multimedia services without compromise, stay with it just a bit longer. It was in understanding the problem of integration that a simple solution was devised.

Integrated and Synchronized Services?

Having summarized the daunting challenge of the physical layer (layer 1 of the ISO OSI model), it is important to address the needs of an end-user application(s) perspective. What is the user looking for now and in the future? Primary and secondary research on the subject of end-user application requirements can be summed up as follows: The goal for integrated-services multimedia applications is quality communications that rival being there while enhancing productivity. For instance, we have witnessed that for design teams, this means being able to conduct a live design-review meeting via videoconference with all parties of the conference simultaneously examining and pointing to and discussing in real-time specific micro elements of a schematic. In this particular case, the conference is optimized by being able to mix static and dynamic content for a truly collaborative communication. For telemedicine, this means being able to conduct a videoconference with patients and colleagues while viewing and describing a particularly difficult procedure in real time. Similarly, the CEO of a large multinational company can conduct a quarterly communications meeting via a multipoint-connected videoconference with satellite locations and third-party partners through which the audience can both interact with the CEO and where the CEO can electronically point to and verbally and visually describe specific cells of a spreadsheet summarizing third-quarter sales, inventory turnover, and the fiscal-year forecasts and have the visual and audio information fully synchronized.

In all cases, the need for integration is obvious. However, what is the value of integrating three forms of communication (voice, audio, and data) if they cannot also be synchronized? Lack of lip synch will have users turn their heads. Lack of data-synch with an audio / video narrative has the potential to misinform and confuse. It is for this reason that it is not enough just to integrate voice, video, and data communications. The integration must yield a synchronized communication that for economic and or convenience reasons makes it better than being there in person. The challenge can be explained by an illustration of a two-dimensional model. *Figure 4* illustrates the different nature of voice and video applications from data. A key QoS measure of merit for voice and videoconferencing is delay. For real-time interactive communications, end users have a very low tolerance for latency for end-to-end round-trip communication. By contrast, data applications are much more concerned with bit-error rate. In this case, the main mission is to move all information (every single bit) from source to destination address without loss or alteration. If an error should occur, data networks are designed to flag and, in some instances, correct or re-transmit a flawed communication. In the case of voice and video, the error rate is less crucial than latency, while delay is less of an issue for data. However, with integrated and synchronized voice, video, and data services, the ideal integration does not result in a compromise of quality for any of the individual communications services.

Quality of Service (a.k.a., QoS) Baseline

Before we address LAN options for integrated services, it is important to consider QoS. Quality must be taken seriously and used as a critical evaluation tool in choosing the medium for bringing even stand-alone videoconferencing capability to enterprise desktop end users. Study after study shows that end users know exactly what they want from a qualitative

FIGURE 4

Simple Human Tolerance Model for Real-Time Interactive Multimedia Communications

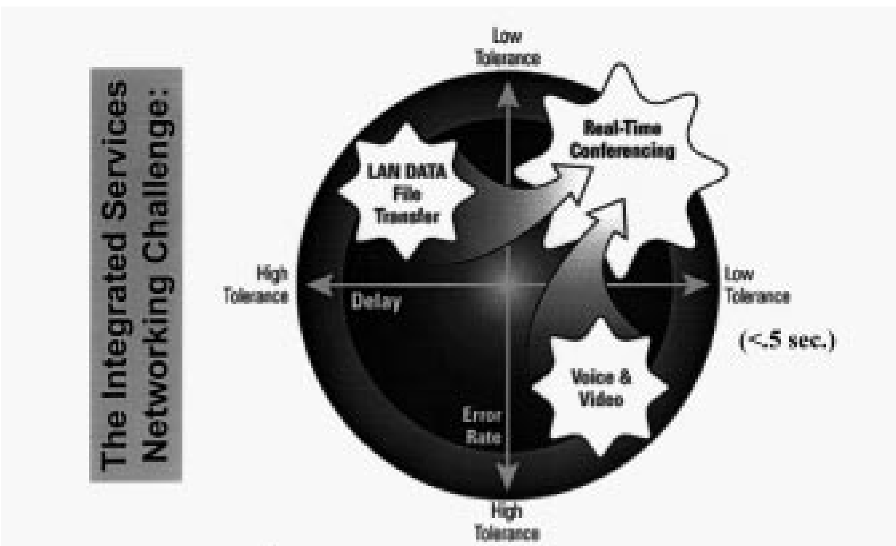

standpoint. However, thus far, the majority of enterprise-network systems fail to meet even the basic needs of those who want networked video applications such as videoconferencing at their desktop. What is worse, a popular publication recently summarized the results of a qualitative analysis of various networked conferencing systems, done by a leading test lab on various networked videoconferencing systems. The labs rated products as delivering "good" and "excellent" performance when they delivered nonlip-synched, 7–9 frames per-second video with audio and video-frame dropout, all on window sizes no more than two to three times the size of postage stamps. There was not even a measurement of end-to-end latency (delay) in the publication's report. Unfortunately, the message from various labs, publications, and OEMs has been interpreted by some end users as "this is as good as it is going to get for now." So, even some of the most ardent leading-edge end users have become disillusioned, causing this particular video application to languish at the enterprise desktop up to now.

As obvious as this may seem, QoS considerations can only be measured on a basis of meeting or exceeding the expectations of the most important critics of all—end users. As we integrate the currently independent communication applications of voice, video, and data, the QoS for each component of newly integrated services must equal or exceed end-user's expectations for each separate application. Expectations for "voice" are based on existing telephone and voice-processing systems, while "data" is based on "the best day" on the LAN / Internet (e.g., during a public holiday with no LAN or WAN traffic, etc.), and "videoconferencing" is referenced to their combined exposure to and experience with $20-$100 room and roll-about systems and broadcast television.

Typically, videoconferencing QoS can be quantitatively measured by image size and resolution, frames-per-second (fps), degree of lip-sync, and the number of micro seconds that are lost due to the time it takes for the capture, compression, LAN / WAN communications, decompression, and rendering of all those real-time interactive images. Research shows that users want fps full-screen images with less than 0.5 seconds of delay, with democratic access and performance parameters, and they do not want to have to learn new user interfaces for old familiar applications. Qualitatively, user satisfaction and resulting user habits greatly affect the overall success or failure of the deployment of desktop video systems.

The biggest challenge from a physical-infrastructure standpoint is to provide for the integration of voice, video, and data services at a quality of service (QoS) level which does not ask the user to compromise on preconditioned (usage-based) expectations for the QoS of each of the independent services. However, there are also issues of quality improvement in the stand-alone services as well (e.g., reliability of LAN data servers and the Internet).

As a precursor to a summary of QoS requirements for integrated-services applications, QoS benchmarks for data, voice, and video communications as stand-alone applications will first be addressed separately.

Networked Computer QoS for Data

In the case of networked PCs, QoS expectations have been based on an age when there was less sharing of the LAN resource due to the combination of fewer users and low-intensity text-based content such as e-mail. The networked PC user has come to expect to be able to send a "one way at a time" (half-duplex) e-mail or print transfer across the network with minimal delay and with a high confidence that the data will not be altered in any way during the transmission. The user expects to be able to send data files to other users on other networks.

With the advent of multimedia and with users now beginning to send large bit-mapped graphics files across LANs and the Internet, the performance of the network has become a growing issue, particularly as it relates to expected QoS. The enterprise data network has become a mission-critical service. Because the Internet and the World Wide Web are rapidly establishing new standards for ease of access to information and global data connectivity, it is also taxing the capabilities of a network infrastructure and protocol which was designed for another purpose and magnitude of use. A prime example of these growing pains was the August 7, 1996, crash of America Online (AOL), the country's biggest on-line computer service. The crash left the company's six million subscribers off-line for 19 hours. Steve Case, AOL's chairman and chief executive officer, based in Chantilly, Virginia, said that the problem involved network-routing software. "We encountered significant problems as we were installing high-capacity switches within the local area network in AOL's data centers," Case said in a letter to members after the service returned. "The real concern is, we have 6 million members who rely on us, and we've let them down. We still have a long way to go to make AOL as reliable as have utilities such as electricity and the telephone."

Notwithstanding the original use of the packet-based LAN and Internet / intranets that connect them, there are numerous companies entering the market which are promoting packet networking as the single platform for all current and future forms of communications. Perhaps the best example of this is found in the fervor surrounding Internet telephony. The concept is simple. Phone calls can and should be made on the Internet because it is free to anyone who already has an networked personal computer on the Internet. For example, whereas a non-discounted telephone call between San Francisco and London might cost the caller a couple of dollars per minute using the traditional phone network, an Internet-based call might only require a $0.04 call to the local Internet service provider (ISP), resulting in a clearly compelling cost argument for packet-routed voice-over-the-Internet phone calls. The same people who back this noble notion of "free voice" over the Internet also back the concept of doing the same for the more bandwidth-intensive applications such as videoconferencing. Given that LAN hub, router, and NIC OEMs are the biggest backers of the concept of packet-routed live multimedia communications, it must be sobering for these companies to listen to 3COM's founder and self-proclaimed "Father of Ethernet," Bob Metcalfe (now of InfoWorld fame), describe his view of voice over the Internet as a "a fraud." He describes his take on Internet telephony in his March 1996 "From The Ether" column, as follows:

> Internet telephony, as currently promoted, will never work. Today's Internet telephony is, well, like encouraging people to heat their homes by burning toilet paper gathered from public restrooms, because it's free! Internet phone calls are intermittent and delayed even during those increasingly rare times the Internet is unloaded, and they're not even close to free.

He goes on to emphasize that all this does not even consider the costly equipment and services Internet telephony users require just to be able to make a "free" Internet phone call, only to get what he describes as a poor-quality service. "You can today, thanks to taxpayers and gas stations, get free toilet paper, but if you were to use much, say for home heating, quickly there wouldn't be any. Calling it free for purposes of home heating is fraud."

Phone Systems QoS

Let us examine end-users' QoS expectations for a system that has been evolving for over 100 years. As a starting point, whether it is a home or enterprise phone, the user expects to be able to pick up the phone handset or press the speaker button at any given point in time and hear a dial tone. The user is programmed for on-demand phone-network access and has a number of other audio-quality and latency expectations which have been programmed over years of usage. One of the most fundamental expectations is that the parties of the call can talk at the same time and interrupt one another. This behavior is the virtual equivalent to sitting at the dinner table debating a point in real time. Imagine the impact on the effectiveness of the conversation if there were a latency (delay) of even a half of a second, or if the only time parties could hear each other is if only one party spoke at a time. It is second nature for us to use the phone as our instrument for instantaneous "on-demand" real-time interactive communication. It is our point of reference for real-time communications.

"Up-time" on phone systems is almost 100 percent. When we engage in a call with one or more parties, we do not even think about the fact that virtually the only event that can terminate that call is for the parties of the call to hang up. The quality of the audio in the call will not drift or diminish simply based on the number of other people who decide to make phone calls. In other words, quality is not usage sensitive.

During the Loma-Prieta earthquake in the San Francisco Bay Area in 1989, the capacity of the phone network was considerably stressed as virtually everyone tried to reach loved ones and friends to get first-hand reports of damage or injuries. It is only in a time of natural disaster that end users come to realize that the phone network is not an infinite mesh. People had to experience network busy signals until other parties concluded their calls. But what it also pointed out is that unless a physical line was broken, all calls that were established before or after the quake were not affected by the fact that virtually everybody was trying to call somebody. The network is designed to guarantee the quality of established calls, and to prevent degradation of quality based on loading. In the end, the expectation is that the sound will not be delayed, broken, or less than a "pin drop" in clarity. Over the past couple of years, users have come to rely on their voice-mail systems to capture and access key information from almost anywhere in the world they, or their calling parties happen to be, at virtually all hours of the day.

Video Systems QoS

There are many forms of networked video end users desire to run over local and wide-area networks (e.g., MPEGI/II, motion JPEG, JPEG, broadcast TV, HDTV, MRI images, H.320 videoconferencing, etc.) but none more intense than those applications which are live, real-time-interactive, and motion intensive. The best example of this challenge is the videoconferencing application. The aggressive and seasoned videoconferencing system user expects to be able to set up a conference with parties within the enterprise and with outside parties at any location worldwide on a whim. Participants in these "conferences" expect to see synchronization between what they are hearing and what they are seeing on the video screen with the ultimate test being the proverbial "lip-sync" test. Furthermore, they expect video with the quality of television and a system which allows for interactivity which rules out any solution which causes end-to-end delay of more than one second from the time a word is spoken and a subtle visual expression can be responded to by users at the other end. All of the other video applications, while potentially requiring lots of bandwidth, do not have to deal with the challenge of capturing, compressing, transmitting, decompressing, and rendering of unpredictable content in real time. An MPEG system can delay the start of a stream of video content a few seconds or longer, and the user will not notice or care. However, a few-second delay in a videoconference call is deadly.

Business Research Group (BRG) conducted research on QoS expectations for networked video applications from the perspective of network planners which yielded some interesting results.[3] BRG's research on the planning expectations of information system (IS) professionals for networked video revealed their expectations for video image and frame rates as illustrated in *Figure 5*. Although many systems on the market today are incapable of true quarter-screen or full-screen images, over 70 percent of the respondents demanded this capability. And even though the targeted respondents were not steeped in video applications, the vast majority demanded a frame rate of at least 15 frames per second.

Perhaps the most telling measure of QoS for a real-time interactive multimedia application was also uncovered by this research. It involved the issue of resource sharing. The respondents were asked what should happen to desktop-videoconferencing quality in the instance of increased demand. As illustrated in *Figure 6*, over 75 percent responded that no degradation in quality should occur to existing calls as a result of new callers on the network. Amazingly, 33 percent of these MIS/IS professionals were savvy enough to suggest that, just like the phone system, a busy signal should be issued until such time that network resources become available to support constant quality of service for all conferences.

FIGURE 5

Simple Video Qos Expectations: Desktop Video Size & Frame Rate

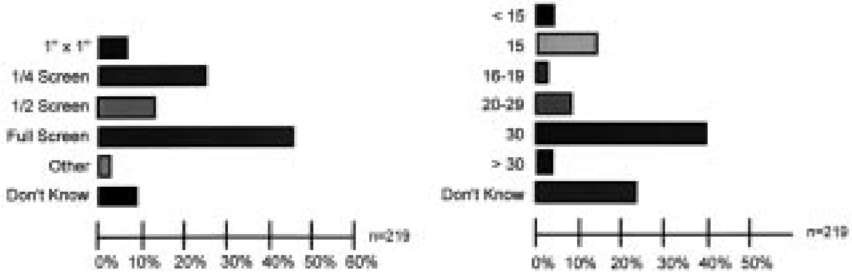

Modeling QoS for Integrated Voice, Video, and Data Systems

As there is a limited historical point of reference for the integrated voice, video, and data application, a tool is needed to predict the level of quality end users will get when they integrate these disparate media types into one system for communications.

The MultiMedia Communications Forum (MMCF) has published an industry QoS document on the subject of networked multimedia applications which breaks down the elements of end-to-end systems in order to model QoS for planning purposes.[4] This model (illustrated in *Figure 7*) divides QoS requirements between terminal equipment, private networks, and public networks effectively creating a performance budget for the whole model and each of its elements. The model specifically defines classes of QoS based on key measurable parameters. The report summarizes that "voice, video, and data components of a multimedia information stream impose stringent requirements on transfer delay (a.k.a., latency), transfer rate (a.k.a., guaranteed bandwidth availability), and error rates in transmission." Minimizing latency is critical for real-time interaction, while guaranteed bandwidth is vital to maintaining a constant QoS for the duration of a call.

For "Premium Performance" (referred to by the MMCF as QoS Class 3), for example, the document details the following multimedia application teleservices requirements in order to achieve 30 frames per second (fps) lip-synched video for real-time interactive conferencing:

FIGURE 6

QoS Expectations for Callers Added

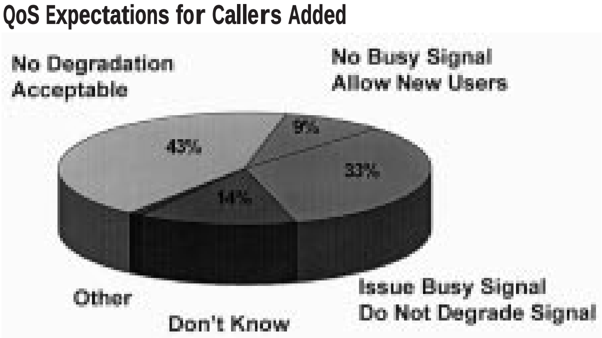

- Audio (voice) transfer delay with echo control of < 150 ms (.15 sec.)
- Video transfer delay of < 250 ms (.25 sec.)
- Video / audio differential (a key measure of lip synchronization) of >-150 ms to <+100 ms)

Clearly, latency in videoconferencing is bad, but how much latency is acceptable to end users? QoS studies by AT&T / Lucent Technologies and other companies have concluded that, while the ideal is zero latency, videoconferencing users are reasonably satisfied with the performance of their systems if round-trip (end-to-end) delay is less than 500 ms. They become annoyed above this level and will not use a service that is over one second in round-trip delay. However, what is commonly misunderstood is that raw network speed does not solve latency problems. Even in calls between San Francisco and London, the latency in the LAN and WAN is measured in the tens of milliseconds, so a doubling or even an order of magnitude increase in LAN or WAN speed has virtually no impact on performance. What matters from a bandwidth perspective is having whatever bandwidth is available on a constant bit-rate and non-shared basis. In reality, the biggest component of delay is the time it takes the sophisticated algorithms used in videoconferencing to perform compression and decompression of massive amounts of content at ratios of 100:1. Systems with good CODERS take 200-250 ms to perform a full cycle of capture, compression, decompression, and image rendering.

Since time is lost by compressing, why should we use compression at all? The answer lies in the nature of how we now communicate. Since we are no longer confining our video conversations to just internal participants, we must use the public network, and, as public networks are tariffed, the cost of broadband network-access equipment and the bill for uncompressed calls would be prohibitive. Therefore, the challenge is to minimize the number of compression and decompression cycles to one for an end-to-end conversation. This may seem obvious, but since the H.320 videoconference standard relies on the telephone network, it, in effect, must have the telephone network's channelization (B-Channels), signaling, and call control from desktop to desktop.

Networking Options for Integrated Services at the Enterprise

What are the apparent network technology front-runners in the race to capture the videoconferencing and broader inte-

FIGURE 7

Quality of Service (QoS) Reference Model

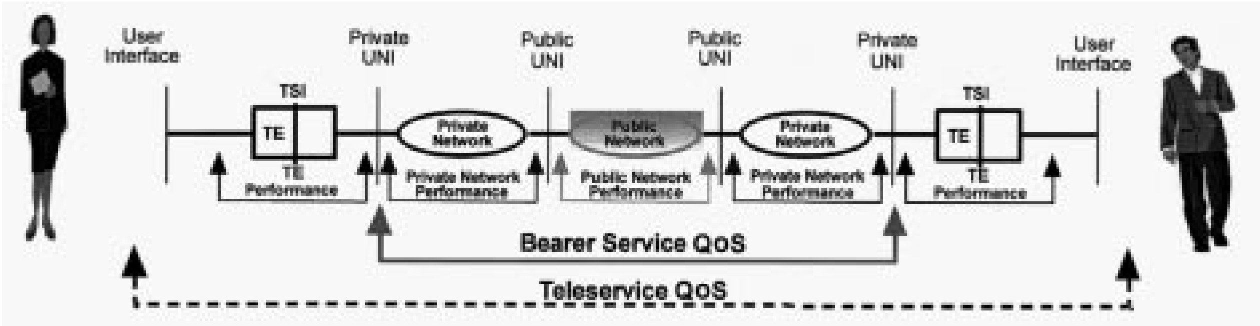

grated-services market? Can any of the current options provide a smooth path to integrated-services multimedia? The options are ATM, packet-based LANs connected to the Internet and intranets, and the new IEEE integrated services LAN (ISLAN16-T) (a.k.a., isoEthernet).

ATM Networks

Although the ultimate potential knight in shining armor, ATM is not the answer for interactive multimedia in the LAN —at least not yet. Unfortunately, until ATM can compete with the QoS and CoS points of the massive installed base of packet LANs for data communications and the circuit-switched WAN for voice, it faces an uphill battle with video as its only potential ally. However, just adding video as another mono-media application is not what users want. End users will not tolerate a compromise of their expectations for existing voice and data systems, and they will not underwrite a forklift upgrade of their existing infrastructure. We have seen a lot of work done on LAN emulation, but it does not look as though the packet-LAN companies plan to concede that packet networks are inferior to cell-based systems for data-networking applications. We have also yet to see a PBX manufacturer advertise the availability of their new line of ATM phones. The good news is that those of us who are ATM backers will keep at it, and just like ISDN, its day will come, and multimedia bliss will permeate the LANd.

Packet Networks

Properly enabled PCs on packet LANs can achieve an incredible CoS advantage over other systems by enabling videoconference calls over the global Internet for the mere cost of a connection to the local IS. Perhaps you have seen a demonstration of some early systems which are designed to do this. If so, unfortunately you know that "cheap" does not compensate for the unmistakable delay and poor quality which results from the use of an inherently non-deterministic and shared medium to handle time-sensitive traffic such as voice and real-time interactive video. Integration and pseudo-synchronization of voice, video, and data applications is an even more random shot in the dark with the Net. It will be this way until truly dedicated network bandwidth can be devoted to video and, for that matter, all forms of communications requiring real-time interactivity. That fix will not be "free." In any event, video systems are being designed to run over popular packet-switched and routed networking solutions (e.g., 10 Mbps, 10/100 Mbps/Fast, and Gigabit Ethernet). However, the integrity of time-sensitive voice, video, and data cannot be guaranteed with packet-LAN-based solutions.

One approach is to dynamically "prioritize" voice, video, and data traffic, one service over the other. Though prioritization schemes are designed to bring low-level congestion relief, there will always be a real concern about reaching an overall traffic threshold that will render the prioritization schemes incapable of maintaining QoS acceptable to end users. Besides, priority carries baggage of its own. What and whose communication gets more priority—the VP making a video call to a prospective customer, or the engineer working collaboratively on an intensive-design database in real-time with an hourly consultant? And this is just the first level of priority encountered within a single company location. Who sets priorities beyond the premises? Ultimately, priority systems mean that someone loses and or pays for priority status. Unless dedicated pipes can be added to today's packet LANs which support the channelization and call-control standards used and required by all of today's phones and videoconferencing systems, adding these services to overtaxed systems originally designed for LAN-based data communications without QoS and CoS compromise is a distant reality. So, is there nothing end users can count on to move into the future of communications without scrapping it all, compromising QoS, and starting afresh?

LANs and Ethernet (a.k.a., IEEE802.3)

When it comes to packet-based networking for LANs, the media-access protocol which is known by the market as "Ethernet" has grown to almost total domination of the enterprise infrastructure. It is a relatively simple protocol which runs on existing "phone grade" cable (a.k.a., Category 3, 2-pair, unshielded twisted pair). Over 100 million PCS are attached to the network this way. Ethernet's media-access protocol is an IEEE standard, and there is a massive OEM, service, and support infrastructure serving the market. The main mission of Ethernet is to provide local-area (premises) connectivity of PCs and workstations for the transference of data (ranging from simple text to large graphic images) from one address to another. All user equipment contains a large and unique binary address which represents the "source and destination" address for all outbound and inbound data communications. This is sometimes known as Internet protocol (IP) addressing. Bridges, routers, and gateways provide the ability for locally connected users to reach and be reached by outside parties.

LANs in general and Ethernet specifically offer a simple value proposition: It works; it is inexpensive to purchase; it is widely supported by OEMs, VARs, and network-support people, and its performance can be scaled through various means to meet the demands of expanded use and applications.

The measure of merit within the LAN business is shipments of network-interface controller cards (a.k.a., NICs) which are installed in the bus of PCs to provide local-area network physical access. Sometimes, these functions are embedded in the system. The industry market-research firm International Data Corporation (IDC) forecasts NIC shipment growth of 12.1 percent annually through 2000 from a base of 29.8 million units shipped in 1995.[5] In 1995, this shipment rate represented a $3.7 billion piece of an industry which makes hubs, routers, servers, network operating systems (NOSs), network management systems, etc. totaling approximately $20 billion in annual revenue to OEMs. By the end of 1996, IDC estimates that slightly more than half (52 percent) of the installed base of 118.3 million LAN-attached PCs will be located in the United States.

Virtually every major analyst forecasts that Ethernet and its derivatives will be the dominant LAN-access method into and beyond the year 2000. Out of a total shipment forecast of 52.9 million NICs in 2000, 44.9 million, or 85 percent, are forecasted to be Ethernet-based. By contrast, another highly touted technology for LAN access, asynchronous transfer mode (ATM) is forecasted to ship 2.5 million NICs in 2000. This is significant as ATM has been marketed as the LAN technology for integrating voice, video, and data systems. Gartner Group's Al Lill forecasts that integrated services will drive the networking market and that in the enterprise, the conduit for these "multimedia" communications will be the LAN, carrying over 75 percent of all traffic before the end of the decade.

Isochronous and Packet Networks

When enterprise end users need to connect to end users at other physical locations beyond their local-area domain, they require access to the wide area network. This network is appropriately referred to as "the cloud" as it is the overall network that connects users of phones, videoconferencing systems, application and content servers, Web sites, and local area networked PCs to one another. This network is also referred to as the public network as it traverses both public and private property. In theory, users of networked PCs and phones can send an e-mail or make a phone call to just about any other person or location in the world through this network.

While this network was originally designed to carry voice communications (analog), it has evolved to where it carrys all forms of information types including voice, data, and video communications. It is truly the information superhighway. The equipment which routes / switches information is primarily based on discrete building-block channels which contain 64,000 binary bits of information per second. These are time slotted with other information on the same medium of connectivity. These blocks of information are known as "bearer" or B-channels. Integrated services digital network (ISDN) is a wide-area network technology / service for delivering these B-Channels to enterprises. There are all forms of ISDN, sometimes referred to as switched digital services (SDS), but the foundation of all ISDN technology is the massive matrix of 64 kbps channels which enable a "connection" to be created between two or more end points. The WAN is built to support a wide range of applications which are ultimately segmented into these 64 kbps channels. Therefore, ISDN has been forecasted to be the dominant WAN-access technology for the vast majority of small to large, enterprise-based end users through and beyond the year 2000.

BRI and PRI ISDN

Two types of ISDN–user interfaces exist: the basic rate interface (BRI) and the much larger bandwidth primary rate interface (PRI). Each of these interfaces enables the line to carry a different number of channels. Within ISDN systems there are two types of channels, bearer (B) channels, which "bear" the data, and signaling data (D) channels, which carry instructions and are coupled with the B channels. B channels are clear connections that are used mainly for dial-up voice, video, and data transmission. D channels have less data-carrying capacity, are shared by the ISDN users on a given physical line, and are used for essential "administration" instructions. The D-channel sends signaling and control information such as call set-up and user-ID packet data. Some carriers support low-bandwidth packet switching on the D channel. Specifically, each U.S. ISDN line with a BRI interface has the capacity, through its switching device and user interfaces, for three channels or pipes for transmission—two B channels, which can carry 64 kbps of data, and a D channel, which allows transmission of 16 kbps of data. Each U.S. ISDN line with a PRI interface has the capacity for twenty-three 64 kbps B channels and one 64 kbps D channel; these are typically used for data transmission through a North American Standard 1.544 Mbps T-1 trunk. Increasingly, companies are using ISDN as a dial-up resource to compliment dedicated "leased line" connections for their WAN-access requirements.

IEEE802.9a ISLAN16-T (a.k.a., "isoEthernet")

The point of demarcation between LANs and WANs in the enterprise has become fairly stark. It is rare to find end users with direct access to SDS WAN circuits at their desktop, even BRI ISDN in the corporate enterprise. End users typically have two cabled networks to their desks: one for their PBX-connected phones and another for their LAN-attached PCS and workstations. The wiring for these two networks is completely separate. The wires are not shared, and rarely is there any linkage between the applications of each.

Although there are customer needs which are best satisfied via a direct connection to the telephony WAN at the enterprise desktop, capital infrastructure and support for doing so is limited, if not non-existent. IsoEthernet is a LAN technology which integrates ISDN channelization and signaling with Ethernet. It is an evolutionary upgrade from ISO/IEC 8802-3 (IEEE 802.3) 10 BASE-T networks.[6] The technology was designed, standardized, and produced by the people of National Semiconductor Corporation. This technology was, in turn, adopted by a group of OEMs and ISVs which manufactured/provide LAN-WAN access systems. IsoEthernet ("iso" for equal, "chronous" for time) is specifically designed to support high-quality store and playback, real-time interactive, and broadcast-video applications in addition to standard-packet LAN applications. The technology extends tariffed narrow-to-wideband telecommunications circuit-switched services (e.g., N*64 kbps PRI ISDN B-Channels) to the enterprise desktop over the LAN cable plant. This kind of network is required for the emerging world of multimedia-rich integrated voice, video, and data applications. As illustrated here, isoEthernet uses the same cable plant and clocking frequency on the LAN (e.g., two pairs of Category 3 or better-grade unshielded twisted pair (UTP) copper wire and 20 MHz clocking) to create a more efficient payload carrier for voice, video, and data traffic. Ethernet, for example, delivers a 50 percent-efficient 10 Mbps pipe while isoEthernet systems deliver an 80 percent-efficient 16 Mbps-payload pipe. Moreover, isoEthernet networks are designed such that voice, video, and data traffic can be transmitted in a bi-directional fashion (a.k.a., full-duplex communications) while Ethernet is a half-duplex medium in over 99 percent of all current installations. Half-duplex communications are reasonable for file transfers but totally unacceptable for real-time interactive (bi-directional) communications such as in voice, video, and even real-time interactive data-collaboration applications. *Figure 8* illustrates how the isoEthernet network connection to the enterprise-desktop time divides 10 Mbps of half- or full-duplex Ethernet packet traffic with 96 discrete 64 kbps (6 Mbps) of ISDN B-channel traffic.

One of the challenges in introducing a range of video applications to the corporate mainstream is in delivering integrated voice, video, and data services at a QoS, performance and price that rivals stand-alone services while leveraging existing infrastructure and applications. Another key factor is to do so without compromising the performance of existing applications. Finally, the ideal scenario is for the end user to be able to use all the services simultaneously without a videoconference call, for example, affecting the performance of a LAN file transfer or a voice call across the isochronous portion for the LAN. Isochronous network services will be the key element of total-system solutions for delivering video quality on globally networked personal computers that rivals the mass market's point of reference for cost and quality for stand-alone resources. IsoEthernet is the only LAN solution which integrates circuit-switched and packet-routed communications resources within the same pipe.

The product solution offered by isoEthernet OEMs is, in effect, an infrastructure which enables the distribution of wideband ISDN in the enterprise over the LAN cable infrastructure. One of the initial horizontal applications uniquely enabled by this infrastructure is standards-and desktop-based real-time interactive conferencing and collaboration, effectively the integration and synchronization of multiple communications media, voice, video, and data. With this integration, the desktop user can now engage in videoconference calls while sharing data files in real time and viewing stored or live broadcast-video content.

The most challenging application in the real-time multimedia arena is H.320 videoconferencing. The International Telecommunication Union (ITU) created the H.320 global industry specification in order to ensure that room and roll-about systems offered by companies such as PictureTel,

FIGURE 8

IEEE802.3 (10Base-T) & IEEE 802.9A (isoEthernet)

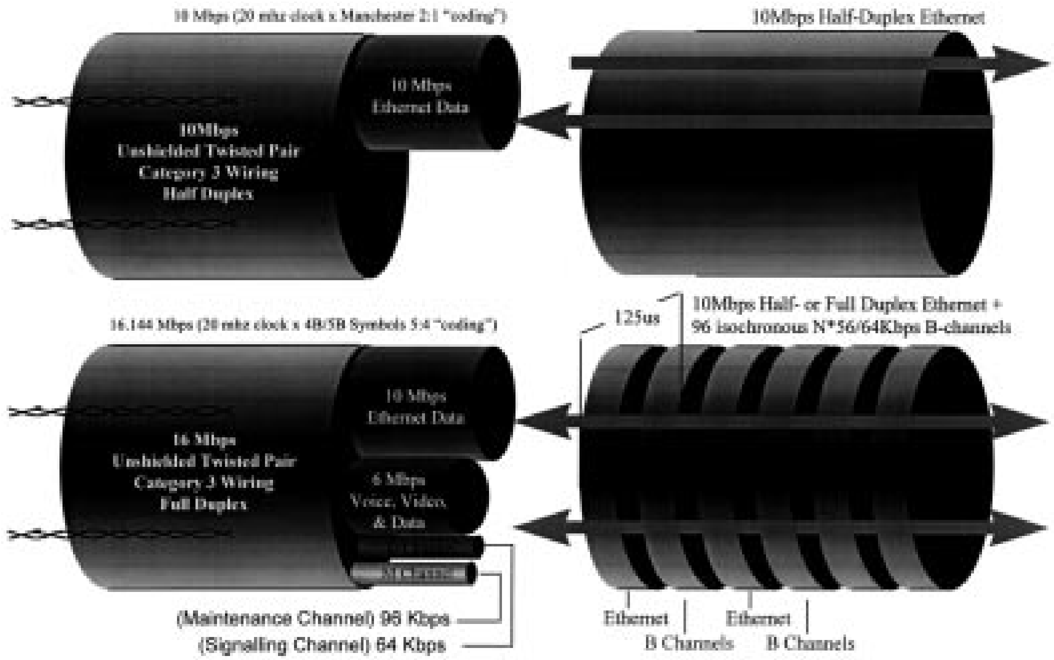

V-TEL, CLI, GPT, Mitsubishi, etc. would interoperate with one another such that end users could be connected to one another from virtually any global end point over the public telecommunications network. The ITU standard is the only open global standard for videoconferencing, and isoEthernet is the only LAN technology which can support the entirety of its requirements across an end-to-end LAN-WAN infrastructure (see *Figure 9*).

One of the specification's unique requirements is for a data-transport framing mechanism which absolutely requires a single synchronization clock for end-to-end connections. The standard also requires that the data pipe between end points be circuit switched, not packet routed. The specification was written around a set of quality goals which, in effect, guarantee the synchronization of audio / voice and video and were / are intended to minimize end-to-end latency (delay) such that parties can interact in a manner which rivals audio communications via telephones while offering an affordable solution. Through the isoEthernet infrastructure, the desktop end user now has access to a solution with video and audio quality that rivals broadcast TV and the telephone respectively. In addition, all existing LAN applications (e-mail, file transfer, Internet access, etc.) are preserved, while key PBX applications are integrated (voice and video mail, hold, conference, transfer, fax, etc.). Vertical applications of these horizontal applications which are enabled by the isoEthernet platform include video banking, distance learning, telemedicine, multimedia call centers, and collaborative design. In each case, the end-user enterprise has begun to use the horizontal capabilities of the integrated services offering for a specific vertical application. An excellent example of this is what Citibank is doing in the area of personal video banking. The company has defined a multiphase program for the deployment of video-banking services. The company has already begun to deploy kiosks in some of its international branches which provide customers with speedy connections to financial subject-matter experts in remote call centers in order to execute a range of transactions from new account and credit-card applications to entire house purchases. In the latter case, customers will be able to connect to loan agents, title-company representatives, real-estate brokers, and other involved parties all in the privacy and efficiency of a single Citibank virtual meeting. Banking can now be done at more convenient hours and with much greater independence of where the involved parties are located.

Deploying isoEthernet Systems into Prevailing Infrastructure

Deploying isoEthernet capability into existing systems is easy and non-disruptive. No changes are required to the installed infrastructure of hubs, switches, NOS, network management, desktop OS, cable and connectors, etc. In fact, because the technology is an IEEE standard that focuses on physical layer 1 which is at the bottom layer of the seven-layered ISO / OSI networking model, all existing systems and software are supported by isoEthernet equipment. For example, this allows support for all existing ISO / IEC 8802-3 collision sense multiple-access collision detect (CSMA / CD) Ethernet media access

FIGURE 9

Bringing the WAN Circuit to the Desktop Over the LAN

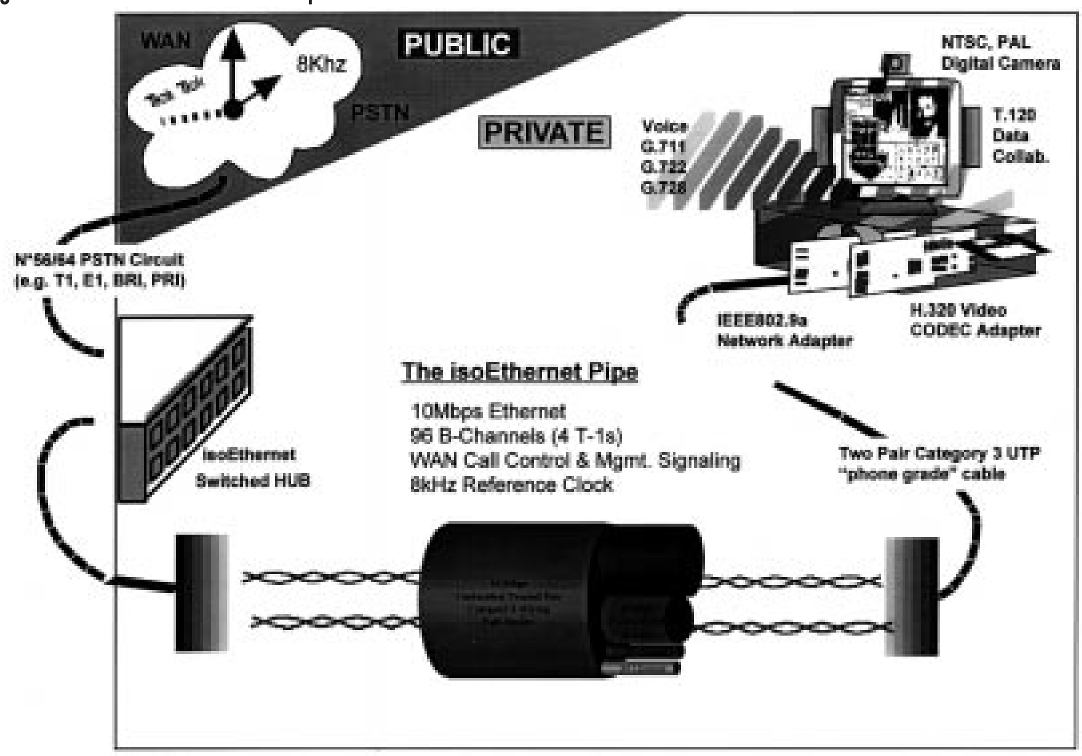

FIGURE 10

An Evolutionary and Incremental Multimedia Application Upgrade

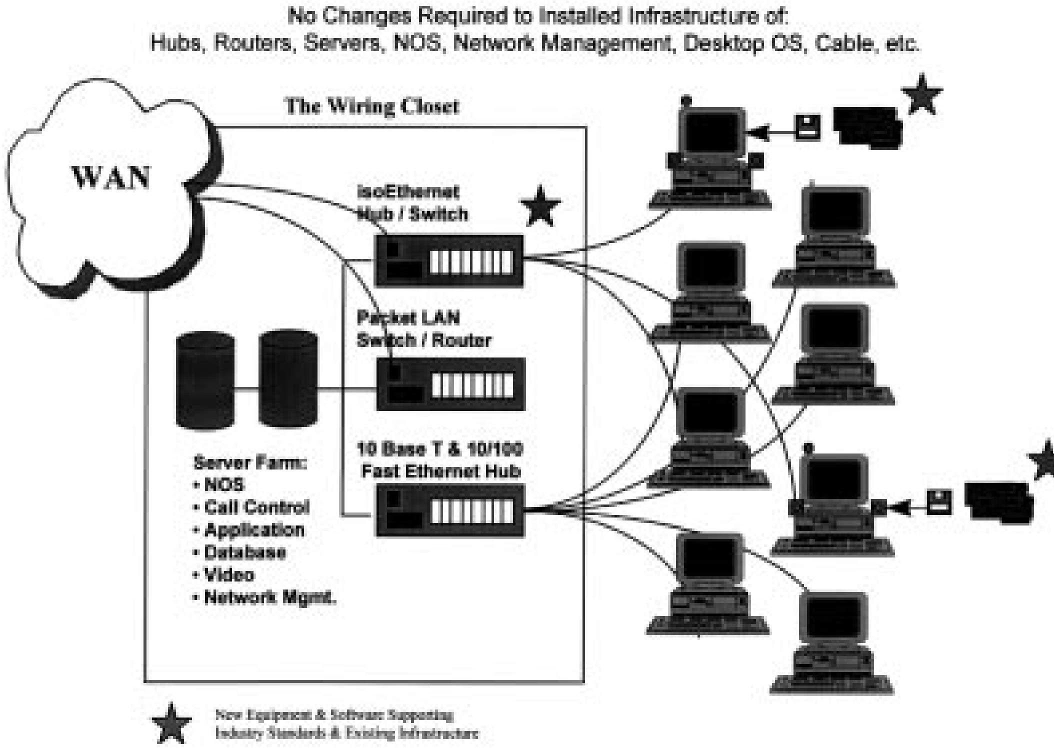

control (MAC) products to utilize the isoEthernet physical-layer interface.[7] In addition, all existing ISDN products (e.g., ITU LAPD, ITU Q.93x, etc.) are also supported on the network.[8] As illustrated in *Figure 10*, an isoEthernet hub can be added to the wiring closet and connected to existing network equipment and services directly.

Desktop systems targeted for upgrade to support interactive multimedia applications will have isoEthernet adapters installed which will replace existing LAN adapters. Once installed, the end user can perform all traditional LAN applications (e.g., file transfer, Internet / intranet / WWW access, network printing, e-mail, etc.) with identical performance and interoperability as with their previous network adapter. However, now the end user will have the discretion of choosing between packet-routed and circuit-switched services to match individual transaction requirements. Furthermore, the end user now has the equivalent of up to four T-1 or three E-1's of switched digital services (SDS) bandwidth at every desktop. As intended by the IEEE 802.9a standard, standard IEEE 802.3 Ethernet adapters can connect to a port on the isoEthernet hub, and any isoEthernet adapter can connect to an IEEE 802.3 hub or switch. The developers were commissioned to develop an upgrade path for LAN systems that would allow for an incremental evolutionary path which would eliminate the need to "forklift" millions of dollars of prevailing infrastructure. Market research conducted by National Semiconductor Corporation identified a distribution of the rate of adoption by end users for integrated-services applications which made it unlikely that the needs of bleeding-edge users would overpower the need to minimize revolutionary upheaval of infrastructure.

Mixed Media Communications and Applications

Any time a new media form is introduced to the end user, a learning experience must take place. In the case of desktop videoconferencing, the typical deployment to date has involved the end user adding a second phone or audio I/O system to the desktop. Ideally, end users would like to avoid the clutter of additional equipment, but more importantly, they are adverse to learning a new set of key punches / strokes for old familiar applications such as hold, conference, and transfer, if these capabilities were even available in their chosen videoconferencing system. Perhaps less obvious, end users who had previous videoconferencing experience are probably accustomed to the room-system experience where, in many cases they did not even need to know how the connection between end points was established. Thanks to the availability of a conferencing administrator / contractor, these were never issues. Everything was planned and orchestrated. This is not the case at the desktop, nor can most organizations afford a personal consultant for every video call end users might want to place. *Figure 11* illustrates how isoEthernet-enabled desktops can address some of these most difficult problems as a result of the technologies ability to interface with both LAN and telecommunications systems in native form. For example, due to the use of industry-standard call-control signaling and channelization, isoEthernet hubs can connect directly in native form to the WAN I/O shelf of most PBXs such that the audio portion of a videoconferencing call can be routed through the PBX. This capability eliminates the end users' need to have a second phone and to learn a whole new user interface for the audio I/O function at the desktop.

FIGURE 11

Distributed PBX and Adjubct Architecture for First & Third Party CTI

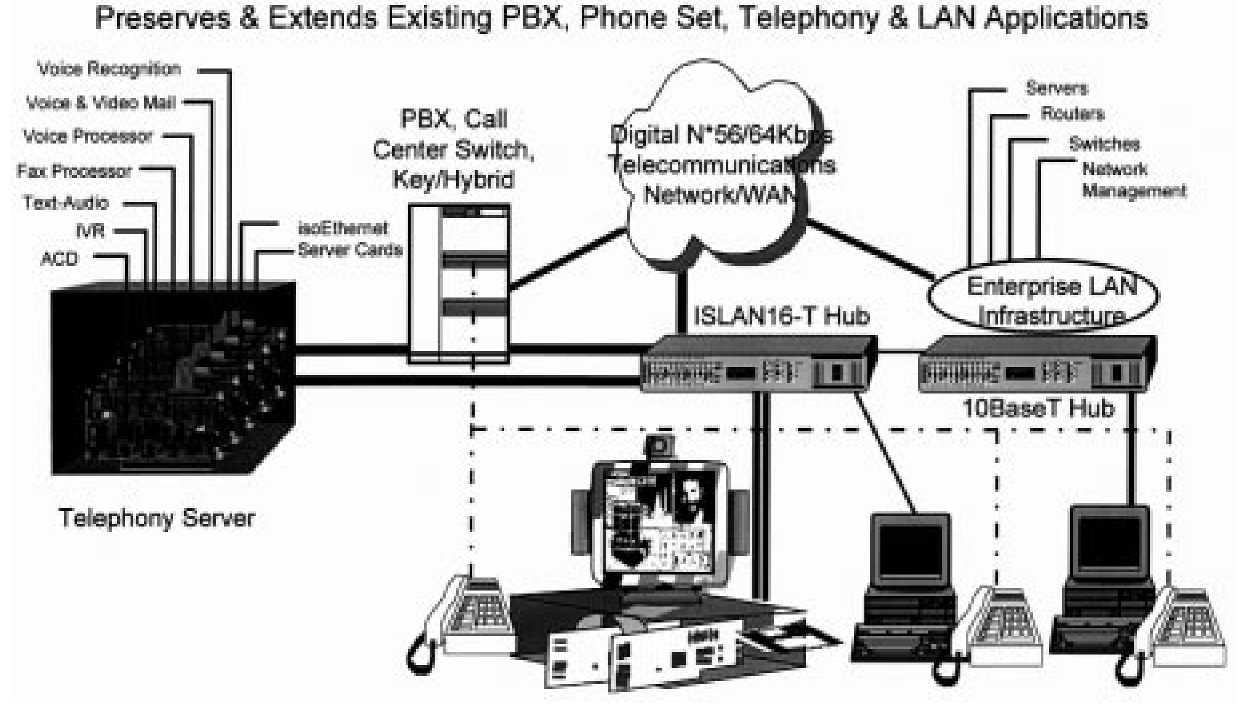

Furthermore, the end user can now integrate the features and applications of their phone systems with the videoconferencing system, a key first step toward true multimedia integration. An example of the benefit of this integration is that the isoEthernet-enabled end users can use the conferencing feature of their PBX phones set to "conference in" parties on POTS phones with other parties which are involved in a multiparty H.320 videoconference call for a true mixed-media communications environment. Another example of the benefit of this integration can be seen when a spontaneous inbound H.320 video call goes to an office with an empty chair. With today's desktop videoconferencing systems, that call might be answered automatically, but no mechanism exists to capture a message. In the isoEthernet environment, after three or four rings, the targeted party's voice-mail system would greet callers with an invitation to leave a voice mail and a promise to call back if they will only leave their number. It does not take great imagination to think about what it could mean to the whole videoconferencing market to be able to leverage traditional voice-processing applications such as interactive voice response (IVR) and automatic call distribution (ACD).

Standardization and Communications Standards

As a result of the customer's need for almost "infinite connectedness," numerous government, commercial industry, and user-group organizations have formed with the expressed purpose of developing "standards" for interoperability. In the case of real-time interactive video communications, industry participants and other parties worked together with the CCITT (ITU) standards body to establish a suite of interfaces to guarantee interoperability between the industry's various videoconferencing systems. This collection of standards is known as the H.320 standard. This particular standard is specifically designed to work on the global digital public switched telephone network (PSTN) and adhere to the standards which the CCITT has established for communications in that domain. Prior to this time, videoconferencing companies devised their own proprietary CODEC algorithms. Similarly, over the past ten years, the IEEE 802 committee has developed a suite of standards to guarantee interoperability of competing vendors' data-communications products in both the LAN and the WAN. All of this has been driven by the user's desire to ensure the ability to interoperate with one another.

ITU Study Group 15

In the interest of defining a basic level of interoperability of H.320 video systems, the ITU's Study Group 15 has defined the mechanisms required to enable users to utilize natively available transports such as LANs, modems, ISDN interfaces, isoEthernet, ATM etc. Work has begun or has been completed for modem / plain old telephone systems referenced as H.324, for non-guaranteed QoS packet LANs (such as Ethernet) referenced as H.323, guaranteed QoS LANs (such as isoEthernet) referenced as H.322, and for ATM, referenced as H.310 / 321.

The following ITU diagrams (*Figures 12* and *13*) and specification excerpts describe the basic elements of the LAN-access approach for supporting the video CODEC portion of the H.320 standard. It should be noted that in all cases, once local information needs to traverse the WAN, it is carried as a digital isochronous SDS service which is billed on a B-Channel per-minute basis. The ITU standard for ATM is not described in detail as all market forecasts show ATM with less than 1 percent of the installed base of desktop connections even by the year 2000.

"The Non-Guaranteed QoS LAN"

This developing standard describes terminals, equipment, and services for multimedia communications over packet-routed LANs, specifically covering the elements needed for interaction with the switched-circuit network focused on support of H.320 video teleservices. H.323-compliant systems must be designed to support the P*64 (H.261) portion of H.320 specifications and must interoperate through gateways with H.320, H.321, H.322, and H.324 networks. The specification does not provide guaranteed QoS. According to the specification's author, "H.323 terminal operation over multiple LAN segments (including the Internet) may result in poor performance."[9] Systems supporting the standard create a "logical" communications path with no relationship to physical path, meaning transmitted information may be lost, duplicated, or received out of order. The standard does not maintain H.320 H.221 framing (8 kHz) to the desktop. Unfortunately, for customers seeking industry standards-based solutions, the ITU specification does not include guidelines or specifications on LANs or transport layers used to connect various LANs, leaving this definition work to the IEEE, IETF, and proprietary implementation.

In order to make connections to H.320-compliant systems, H.323 systems require a gateway for LAN packet to SCN circuit conversion and "gatekeeper" to perform IP addressing to ITU E.164 phone-number identification translation. It is anticipated that these conversion functions will cause serious performance problems which will be reflected in poor QoS for end users. To leverage off and improve on the Internet protocol point-to-point "best effort" protocol, the IETF and competing OEMs are working on methods to provide traffic-flow resource management, routing, and priority to improve quality and add features. New standards are also being defined within the IEEE 802.1 (bridging and switching), and 802.3 (full-duplex Ethernet) to support elements not covered by the H.323 specification.

ITU-T SG-15 H.322 "Guaranteed QoS LAN"

This standard describes the terminals, equipment, and services for multimedia communications over LANs which are configured and managed to provide a guaranteed QoS equivalent to that of N-ISDN (up to 2 Mbps) such that no additional protection or recovery mechanisms beyond those mandated by Recommendation H.320 need be provided in the terminals. Any solution claiming to be H.322 compliant must support not only the P*64 (H.261) video-compression algorithm portion of the H.320 specification but the rest of the full H.320 standard. For example, the standard requires that the ISDN clock (8 kHz) is available at the terminal and H.221 framing structure for a 64 to 1920 kbps channel in audio-visual services. H.322-compliant solutions are more generally applicable to connecting N-ISDN terminals than just visual telephone systems (i.e., are not dependent on the content of the signals carried). H.322 also requires direct support for the E.164 phone-numbering plan for the ISDN era (same as I.131). H.322 therefore does not involve the Internet protocol (IP)-to-E.164 translation "gatekeeper" function required in H.323 systems. An example of a suitable LAN is the IEEE 802.9a "isoEthernet" Integrated Services (IS) LAN16T which provides isochronous services with CSMA / CD MAC service. It should be noted that while the ITU uses the same diagram for all H.32X standards, the H.322 gateway block shown in this diagram is integrated in the central hub of what is the star wired 10 BaseT topology and is therefore capable of connecting directly to the N*64 kbps ISDN WAN without an intermediate gateway or gatekeeper function. Because isoEthernet also supports the Ethernet MAC, it can simultaneously deliver all the services available through H.323 and from the work of the IETF, the IEEE 802.1 (bridging and switching), and 802.3 (full-duplex Ethernet) as they become standardized and available. IsoEthernet-based systems have this added flexibility but are not dependent upon these new standards, as is the case with all H.323 packet-LAN solutions.

Integrated and Synchronized Voice, Video, Data Services Conclusion

The challenge in providing end users with truly integrated and synchronized voice, video, and data systems is that end users are now starting to see a blur between expectations for these heretofore "stand-alone" services. They have come to expect that "the technologists" will create solutions that allow them to have integrated services without compromise. They further expect that, with integration, they will get a higher aggregate QoS than if the services were left independent. And they are right. For example, in the case of videoconferencing systems today, there are three main drawbacks for users: you must schedule your call in advance; you must leave your office to engage in a call; you will not be able to integrate other forms of personal data and content in the call. With LANs, users are being limited in their exchange of certain bandwidth-intensive multimedia forms of content (e.g., video clips) across the network, because they must share the network with traditional "mission-critical" e-mail, file-transfer, and print-sharing applications. Phone users today are limited in their ability to see the physical expressions and responses of those they are talking to. These are just a few of the current day limitations that users encounter, and some "early adopters" believe will go away with integrated desktop multimedia services.

An industry-standard networking technology has been created which blends circuit-switched ISDN channels with packet-routed Ethernet known as isoEthernet. It adds a dedicated "extra" pipe to the packet network, which extends up to four T-1 or three E-1 spans worth of ISDN bandwidth from the circuit-switched WAN over the LAN to every enabled desktop. It seamlessly integrates networked computer and telephony system infrastructures, while operating transparently with installed Ethernet networks totaling over 100 million clients and growing. New "stackable" hubs are added to the wiring closet while desktop PCs are selectively targeted for new adapter cards and videoconferencing system add-ons. Nothing is discarded or displaced—same NOS, OS, dri-

FIGURE 12

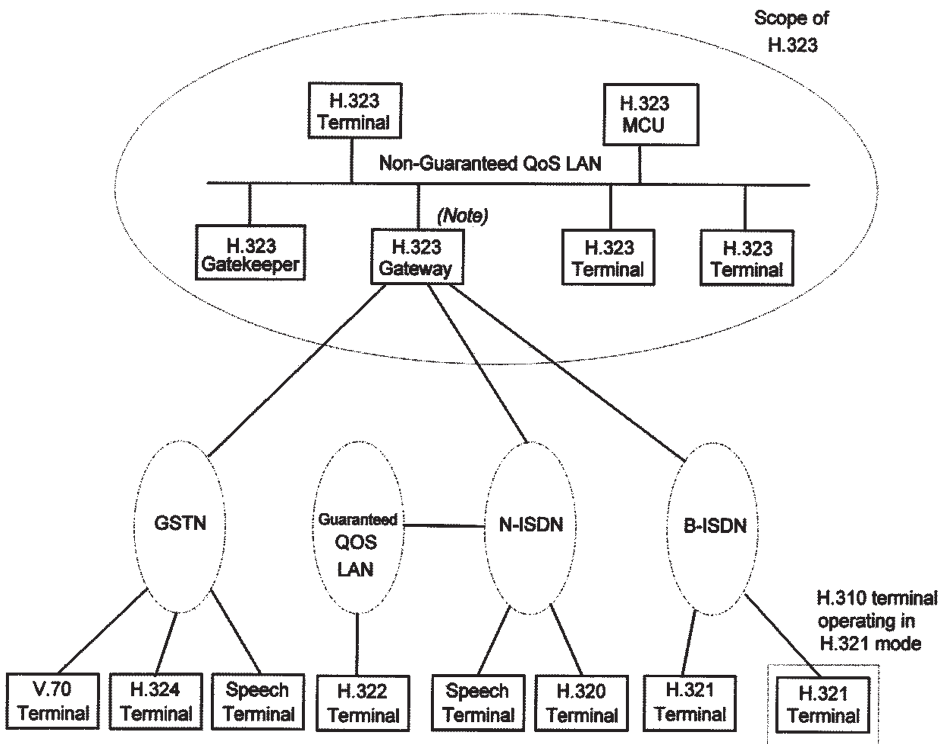

FIGURE 13

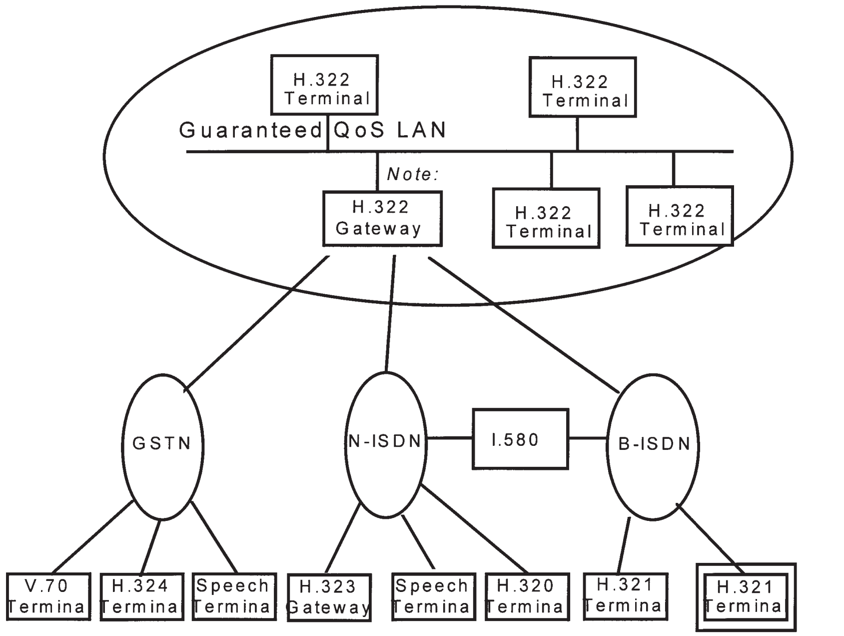

Note: The H.322 gateway in IEEE802.9a i(isoEthernet) s integrated in the central hub of what is the star wired 10BaseT topology and is therefore capable of connecting directly to the N*64 Kb/s ISDN WAN without an intermediate gateway or gatekeeper vers, servers, cables, routers, network-management software, IP / IPX applications, etc. More importantly, a new capability has been deployed which provides a vital network platform for integrated voice, video, and data services. First-time users cannot believe that they can conduct room-quality multipoint video calls with linked and synchronized data applications to users all over the world, while simultaneously playing MPEG clips from the local server, watching live broadcast TV, surfing the Web, and sending e-mails and files to any IP / IPX address of their choice. What is perhaps most impressive to end users is that they see no impact to the performance of any single application as one application after another is transported on the same network. This is the end-user impact of non-shared, non-prioritized communications. From a videoconferencing standpoint, isoEthernet is the only network technology that provides a fully compliant transport for CODECs supporting the full ITU H.320 specification. From the end-user perspective, the ideal world is a non-proprietary heterogeneous place, where our systems enable us to communicate to anyone, anywhere, anytime without concern for the standards that make it all possible. As for its quality, isoEthernet is the only example referred to by the ITU as the "Guaranteed QoS LAN."

References

1 Robert Mirani, Senior Telecommunications Analyst, The Yankee Group, 200 Portland Street, Boston, MA 02114, DVC Conference & Exposition, September 5, 1996.

2 "AOL goes AWOL All Day," Associated Press, San Jose Mercury News, August 7, 1996.

3 Ball, Cheryl, "Market Opportunities for Networked Video Business Research Group," Newton, Massachusetts, 3 / 95.

4 "Quality of Service Reference Model" MultiMedia Communications Forum Document, ARCH / QOS / 94-001, Rev. 2.0, Sect 6-10, 1996.

5 Smith, Justin, "Local Area Networks," 1996 PC NIC Market Forecast, 1995-2000, Filing #11474, Vol. 1, International Data Corporation, Farmingham, Massachusetts.

6 "IEEE 802.9 Isochronous Services with Carrier Sense Multiple Access with Collision Detection (CSMA / CD) Media Access Control (MAC) Service," IEEE 802.9a Draft D10, March 30, 1995.

7 "Carrier Sense Multiple Access with Collision Detection (CSMA / CD) Access Method and Physical Layer Specification," ISO / IEC 8802-3, 1993.

8 "ISDN User-Network Interface Network Layer (Layer 3)," CCITT (ITU) Recommendations Q.930-Q.940, 1988.

9 Quote and Diagram Source: ITU-T SG 15 Document AVC-870R3.

The Role of Electronic Data Interchange in the Evolution of Electronic Commerce

Hilary B. Thomas

Interactive Telecommunications Services Inc. (ITS)

This paper addresses the convergence of two different aspects of electronic commerce. It is abstracted and updated from a detailed study conducted by ITS in late 1996.[1]

EDI + EC = ECI

Two separate and powerful forces are converging from different directions on the electronic commerce marketplace. Electronic data interchange (EDI) has developed to address the need for intercompany systems to interact across company boundaries and between trading partners. Electronic commerce (EC) has emerged as a process-oriented set of functions that enable business transactions, including payments, to be completed between enterprises, their trading partners and customers. From the consumer interactive services perspective, electronic commerce has evolved to encompass transaction-based services such as order entry, purchasing, shopping, and banking, bill payment, and other financial services. These two aspects of electronic commerce are in the process of evolving, separately and together, towards a common place that is being called "ECI." The concept of ECI is outlined in *Figures 1* and *2*.

Electronic Data Interchange

The role of EDI is to take documents which have been prepared manually or by automated systems (e.g., inventory systems), to translate those documents into a standard electronic format and transmit them to trading partners outside the boundary of the enterprise. At the receiving end, the standardized, transmitted form (the "transaction") is "untranslated" into a format that can be read by the recipient's automated application. The result provides the potential to eliminate the delays, errors, and inefficiencies inherent in paper-based transactions. While EDI initially focused on the automation of procurement and order entry systems, it increasingly integrates additional processes in the supply chain in a move towards end-to-end automation.

This paper was first published by the Interactive Services Association in the "ISA Strategic Briefing Book Vol. II," and is reprinted with permission.

The success of EDI has been based on the ability of the industry to develop and agree on standards and protocols for these exchanges. The universal EDI standards lay out acceptable fields on a business form. Every EDI message must specify and use these standard business form structures so that the information can be easily translated upon arrival. There are two competing EDI standards for value-added network (VAN)-transported EDI and pairs of trading partners use one of the following:

- EDIFACT business form standards, which were developed by the United Nations Economic Commission for Europe and are based on TRADECOMS out of the UK. EDIFACT is becoming the de facto standard for international EDI (although ANSI X12 is popular in the Far East).

- X12 business form standards were developed by the American National Standards Institute (ANSI). The ANSI X12 standard is widely used in the USA. It maps a traditional paper document into an electronic format that can move easily over telecommunications networks. Various portions of the standard actually describe individual "forms" for use in specific transactions.

Most EDI software development companies are dedicated to standards although some are beginning to develop alternatives and to create transactions between the two. As with any "standards war" the reality of two incompatible EDI standards is not helping the wider utilization of EDI in domestic and international trade. Efforts are under way to unify the two, but it is unclear whether a straightforward and cost effective migration of existing systems toward compatibility is possible or meaningful.

For the U.S. Government and for some companies, EDI is a requirement for doing business with key trading partners. Large manufacturers and retailers—automotive manufacturers and grocery chains, for example—often require that their suppliers use EDI as a precondition for all purchase agreements. These large "hub" sites are in a position to dictate the use of EDI to their smaller trading partners, or "spokes." However, EDI has been used effectively only among large players.

In 1995, the biggest EDI hub companies averaged more than 400 EDI spokes; this represents between 40 percent and 50 percent of the companies with which the hubs do business. The other 50–60 percent are often smaller trading partners with which the hub has a less well-established relationship,

FIGURE 1

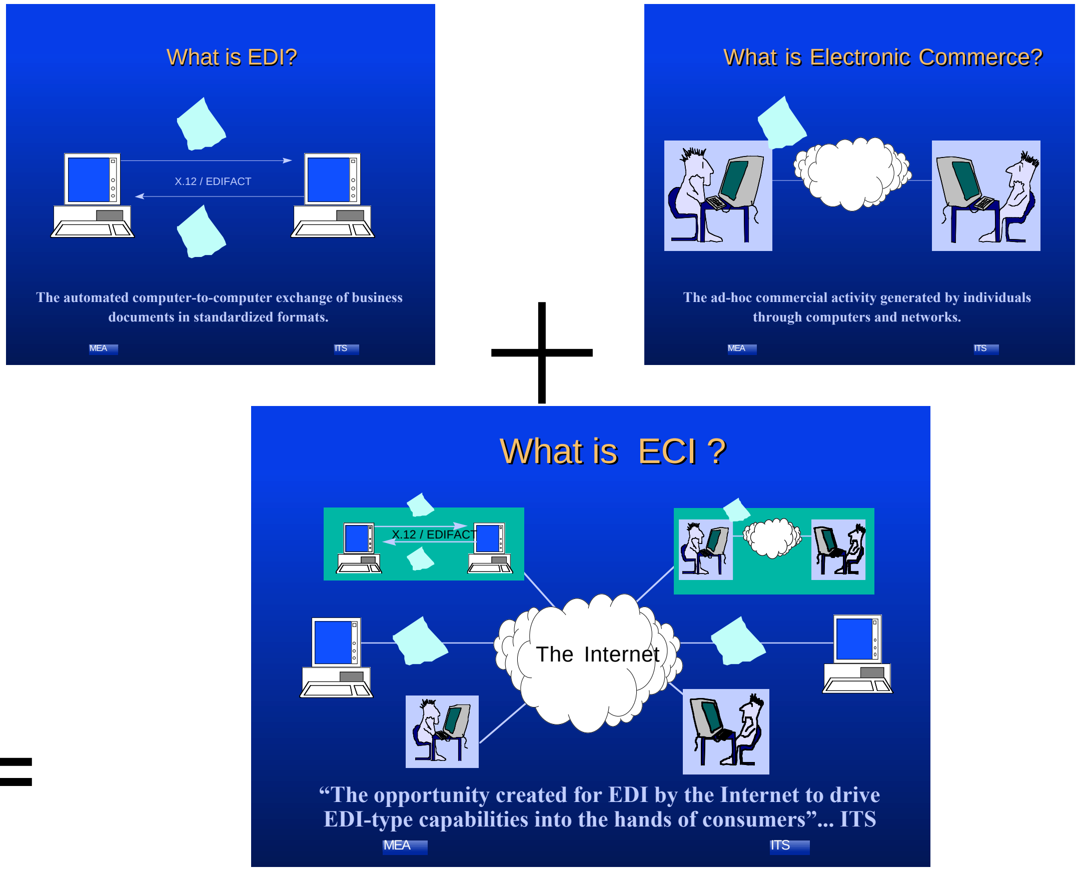

FIGURE 2

The Evolution of Electronic Commerce

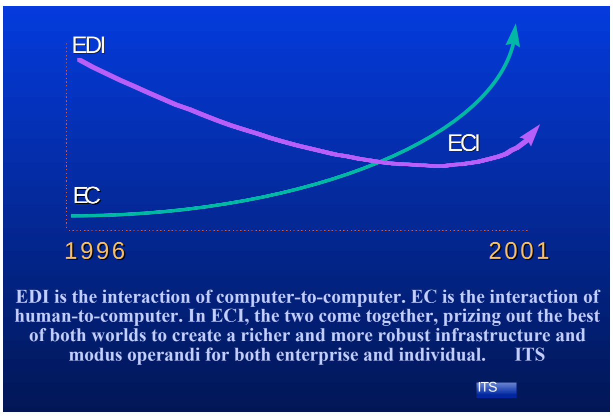

partly because of the lack of EDI connections. Approximately 100,000 companies worldwide use EDI; this number represents only five percent of all businesses that could benefit from the technology, according to Forrester Research.

Electronic Commerce

The term electronic commerce has been used to describe any business-to-business or business-to-consumer transaction or purchase where one or both parties use a computer, terminal, or automated application to place, follow up on, or effect delivery of, an order. Electronic commerce and EDI are sometimes incorrectly equated. There are several distinctions that should be made between them.

- Electronic commerce covers the domain of EDI and much more.
- A fundamental goal of EDI is to improve, by automation, the interchange of information between suppliers and partners, particularly in commercial and government trading.
- By some definitions, EDI only specifies a language or format for business information.

For purposes of this discussion, electronic commerce encompasses a wider range of commercial applications between trading partners and customers, which includes purchasing, ordering, customer support, and the other functions of EDI as well as other business-to-business interactive services and interactive services to consumers. These may include on-line ordering of goods/electronic shopping, bill payment, and online banking, which may be accessed by users located in a place of business or in the home.

The focus of EDI has been automation of business processes and the reduction or elimination of human interaction. The focus of EC has been reaching out to the customer, encouraging simpler and lower cost access to computer-based services by individual customers.

EDI on the Internet

It is no surprise that companies around the globe are looking for ways to increase access to EDI-type capabilities. Nor is it a surprise that the industry is looking to the Internet to see if it offers a means to bring the other 95% of potential users into the fold. EDI transactions have traditionally been transported via third-party value-added networks (VANs). It has not take the industry long to realize that the Internet has the potential to compete with or to augment these VANs, to provide a ubiquitous, less costly, and easier method of interaction and far greater reach.

In an effort to take advantage of the extended capacity promised by a standardized global network, thousands of companies have set up Web sites or have activated some level of presence on the World Wide Web. These companies are attempting to disseminate information, stimulate consumer interaction, and to use the Web as an advertising medium or for customer service. The Internet is predicted to change the way in which consumers handle many commercial transactions and a frenzy of activity surrounds attempts to remove critical barriers such as the high perceived security risk and lack of appropriate payment systems. At the same time, an increasing number of organizations are seeking to offer business-to-business services to a wider universe of trading partners exploiting the universal accessibility and standardization of the Internet. The current opportunity to broaden EDI capabilities to the Internet rests upon this desire to expand commerce more cost effectively to more trading partners.

Key to this discussion, and at the core of the business opportunity, are the following:

- the use of the public Internet versus, or in addition to, VANs to transport EDI and EDI-type commercial messages efficiently and at low cost to a universal audience

- the creation of a standards which will drive development of the infrastructure and implementation of applications based on the new protocols

An Historical Perspective

While there is much debate about the extent and speed of change that will be caused by the Internet, there is no doubt that it is a significant force and that it has already had substantial impact. The opening of the Internet to commercial and individual use represents a shift from "closed" or “private" network computing to "open" access, both for the consumer and for businesses. Interactive services have been evolving for over twenty years but until the Internet era they were largely dependent upon closed proprietary systems and limited business and/or consumer audiences. The client/server model makes open network computing available to users in business. The promise of open network computing is value for users in every segment and strata, both among businesses and in the mass market of consumer commerce. Indeed, as is demonstrated by the French Minitel, one significant element of the success will be the blurring of the distinction between "consumer" and "business" use.

Some companies take the view that the Internet has the potential to revolutionize many aspects of the way we live our lives and go about our business. However, while technology may have the intrinsic ability to revolutionize, people have an intrinsic tendency to “evolve.” This evolution embodies the innate tendency which people and organizations have to adopt technology slowly, over years or even decades. A pragmatic adoption model, applied specifically to electronic information services, appears in the developmental history of these services. It appears that it is emerging in relation to the Internet and that it can be illustrated as follows:

Stage One: Innovation and Technocracy

A technology is invented and "looks for something to do.” Technologists claim an imminent revolution. Academics and entrepreneurs start to develop innovative products. High-tech companies invest in research and development and early adopters use the products and services.

Stage Two: Advertising and Lost Shirts

Nobody makes any money. Industry analysts say that revenue will come from consumer advertising but very little money comes from advertising. Companies pull out, having depleted budgets with little short-term returns to show.

Stage Three: Business Before Pleasure

Companies and developers begin to see opportunities for business services. Industry specific, commercial and transactional services begin to make money. Investment in business services creates the infrastructure and resources for consumer transaction services.

Stage Four: The Mass Market

Companies leverage the original infrastructure and adopt new techniques to reach consumers in ever-larger numbers. Money starts to flow into advertising and consumers communicate and transact.

Today the Internet is moving from Stage Two into Stage Three. More companies are involved and fewer have pulled out than in earlier technologies such as videotex, probably because of the visibility of these efforts and the expectation of low cost of entry. There is no doubt, however, that millions of dollars, if not billions, are being spent by American (and global) businesses building Internet presence without a very clear idea of how that presence is going to impact core businesses, revenues or costs. Just like the original focus of videotex, the focus of the first years of the Internet era has been reaching the consumer. The catalyst was the World Wide Web and other developments that made access to the Internet feasible and easier for “ordinary” users. The Web has become synonymous with the Information superhighway and governments and pundits alike proclaim that the answer to the information age is at hand.

In 1996 it was recognized that the greatest impact of the Internet, at least in the next three to five years, might be in the realm of business-to-business applications rather than the mass consumer market. “Intranet, intranet, intranet,” became the mantra of Microsoft, Netscape, and others staking a serious claim in the revolution. It is now possible for many to find a reason for investing in the Internet. Many companies are looking to the Internet and intranets to improve access to electronic information and services, to increase effective communications within and between companies, and to offer more direct service to trading partners and customers. The focus is increasingly on ECI.

The Parallel Evolution of EDI and Electronic Commerce

In the "WWW.era" (the world since 1994) interest in Electronic Commerce has escalated beyond all expectations. This interest has evolved from two directions:

1. intra-company [or enterprise] process-management
2. inter-company and consumer interactive services

Enterprise process management has evolved from the monolithic, mainframe-based systems of the 1970's to the process-centric, personal-computer / desktop-based systems of today. Client / server systems, graphical user interfaces (GUIs), object-based distributed systems, and increasingly wider access from local area networks (LANs) to wide area networks (WANs) to VANs and now to the Internet, have opened up processes within companies and enabled integration across functions and departments. Products based on EDI's X.12 standards and others have both contributed to and benefited from this evolution.

The interactive services “industry,” which emerged from the consumer teletext and videotex services of the 1970's, has been involved in the development, testing, and marketing of both business-to-business and business-to-consumer services since the early 1980's. Even this, though, never saw anything like the Web and Internet activity of the past 24–36 months.

Initiated by the European PTTs in the 1970s, teletext and videotex enabled the first attempts at "universal" access to computer-based information services—initially for consumers and later (in the early 1980s) for inter-company applications. Along with the development of client / server, terminal-based networked MIS, this represented the first move towards i) universal access within companies and ii) inter-company processing.

During the latter half of the 1980's the focus of interactive applications shifted to consumer online services. Initial attempts were made to launch advertising based information services such as electronic newspapers, e.g. Knight-Ridder’s Viewtron service in South Florida, without much success. Later came the focus on home shopping, home banking, and other transaction-based services. Terminal-based solutions were almost all complete failures. Small successes began to come with the proliferation of home-based personal computers with modems and the emergence of competing on-line services such as America Online and Prodigy. Of the on-line services, only CompuServe successfully targeted both businesses and consumers. Now that the World Wide Web has simplified and codified access to on-line services, both businesses and consumers are potential targets for electronic commerce.

Advantages and Disadvantages of the Internet for ECI

At first glance the Internet would seem to offer the perfect solution for electronic commerce, both as means to extend EDI to new users. It has the potential to be the vehicle by which consumer interactive services (shopping, banking, etc.) will finally become a reality. An investigation of investment and the rate of innovation related to the Internet suggests that ultimately it will become the carrier of choice for the bulk of the world’s electronic commerce traffic. It certainly offers several advantages as a medium for electronic commerce:

- It is, increasingly, easily accessible globally.
- It is based on an accepted, though evolving, set of standards.
- It offers flat, volume-independent, and time-of-day independent pricing for data transmission.
- It is robust because it encompasses multiple alternative pathways, gateways, and interconnections.
- It has high bandwidth for data throughput and it is platform independent.

EDI will move to the Internet. In the short term, cost pressures will motivate the industry to look for a cheaper transport mechanism. In addition, the opportunity to serve a wider audience will move volume off VANs and to the Internet and / or encourage VAN operators to adopt more “Internet-like” approaches both technically (standards) and commercially (pricing). In the longer term, there will be an explosion of investment by both start-ups and mature companies in Internet technology, in intranets, in applications based on the Internet protocols, in security, accessibility, agents and avatars, currency, and other cyber-commodities yet to be imagined. This will create new standards and a global infrastructure, which will eventually supplant or overlay the traditional EDI infrastructure and offer a radically different environment for electronic commerce. As with any new medium, however, there are significant concerns, regarding the suitability of the Internet for “mission-critical” commercial transactions.

Most important among these concerns are:

- perceived lack of security
- inability to confirm message integrity
- vulnerability of messages to interception and fabrication
- lack of user support
- difficulties in obtaining reliable assurance of authenticity or receipt
- international trade concerns about e-currency and other issues

These are the issues that have slowed efforts to proliferate business transactions over the Internet. In spite of the inevitability of its “success,” the Internet is not yet a trusted medium and is open to question and debate over its suitability for business-critical transactions.

John Berry of GEIS says " the Internet allows me to reach more customers more cheaply but it does not replace the real value-added I offer on my private network. After all, it's only another network."

Within five years, however, most EDI will transact over the Internet, predicts Jack Shaw, president of Electronic Commerce Strategies, Inc. in Marietta, Ga. Conventional EDI users are weighing their options. In a recent survey by Forrester Research, Inc. in Cambridge, MA, 20 of 30 heavy EDI users said they are looking seriously at using the Internet in conjunction with the existing EDI infrastructure.

Summary of Benefits of the Internet for EDI

In summary, four factors make the Internet extremely attractive as part of a company's EDI infrastructure:

1. *Flat Pricing*. The standard VAN pricing model incorporates charges per character while users may benefit by using the Internet's current flat pricing model, which is not dependent on the amount of information transferred.

2. *Cheap Access*. Business users normally have a relatively low connection costs, typically a simple flat monthly fee for leased-line or dial-up access. There are numerous commercial and non-commercial Internet access providers in nearly 140 countries. Network coverage and access is substantial.

3. *Common Mail Standards*. More traditional EDI services rely on proprietary software and front ends. This limits the ability to operate with other possible users and narrowed application choices. The Internet has common mail standards, proven networking, and inter-operable systems. These mail standards are nonproprietary and handle congestion and message handling extremely well. VAN network e-mail messages can take hours or days to reach a final destination, while it may take only hours or minutes on the Internet.

4. *Security*. Security is a major concern that is being addressed by the incorporation of public key encryption methodologies into e-mail systems and by the development of "private-public" network features. These methodologies will ensure the privacy of EDI messages and give users a way to verify the sender or recipient.

Implemented via the Internet, EDI-type systems, which are based on proprietary systems and VANs, will bring further benefits to a greater number of companies and markets. The wide acceptance of the Internet will enable a broader range of businesses to utilize networking technology for their business dealings. This includes small- and medium-size companies, as opposed to the past when only large corporations could afford investment in EDI technologies. It also provides a cheaper transport facility to companies who have already been involved in EDI and who might have been using expensive dedicated resources.

In addition to the increased types of business, the Internet can provide more flexibility in terms of the following:

- locating suitable and competent business partners
- establishing short-term and ad hoc business relationships between enterprises
- opening up these features to consumers

Perhaps the greatest opportunity for EDI presented by the Internet is the potential to drive EDI-type capabilities into the hands of consumers. These capabilities include agent-driven applications that, without human intervention, will make purchases and perform other commercial functions for individuals in a way that not even large businesses have used in the past.

The Evolution of Electronic Commerce

Figure 3 illustrates the four phases of evolution of electronic commerce. During the first two phases (1996 - 1998) traditional EDI remains at the core of electronic business-to-business transactions for larger companies and for "mission critical" events. However, it will be increasingly influenced by the new standards with users starting to access EDI hubs, indirectly at first, and later more directly, through translation products some of which are already available in the market. In the second phase, there will be increasing emphasis on the creation of user-friendly interfaces for access to these systems by "humans" as well as by computers and the focus will change to "masking" of the EDI standards. These will still be the primary means of transport of computer-to computer transactions even though the original entry point of the transaction may have been an individual at a desktop connected via the Internet. The automated system will develop to have a human "face."

Initially there will be a significant distinction between the "external" Internet transaction and the "internal" flow of data over the more secure EDI VANs, which are at the heart of EDI. However, moving forward towards Phase Three, the number of new hubs using the old EDI standards will decrease and the role of the VANs will to shrink in relation to the whole. The amount of traffic will increase as more, smaller trading partners access EC/EDI via secured intranets (the private use of the public Internetwork), and as the translation and masking of EDI gives way to new standards for electronic transactions between companies based on TCP/IP and the Internet.

Phase Four represents the "Great Change." It will probably not arrive with a big bang. It will evolve through the first three phases but, by the year 2000 (but not before), it is probable that the predominant milieu for automated transactions between companies will be Internet/TCP IP based. VANs will not have disappeared but they will support TCP/IP protocols in order to survive; no more new EDI Hubs will be set up based on today's EDI standards. By this time, it is also anticipated that consumer electronic commerce will be emerging as a force in the way in which trading takes place between retailers, banks, utilities, and other suppliers of goods and services and their consumer customers.

Over the next five years we will see increasing convergence between the way in which businesses conduct business electroni-

FIGURE 3

The Evolution of Electronic Commerce

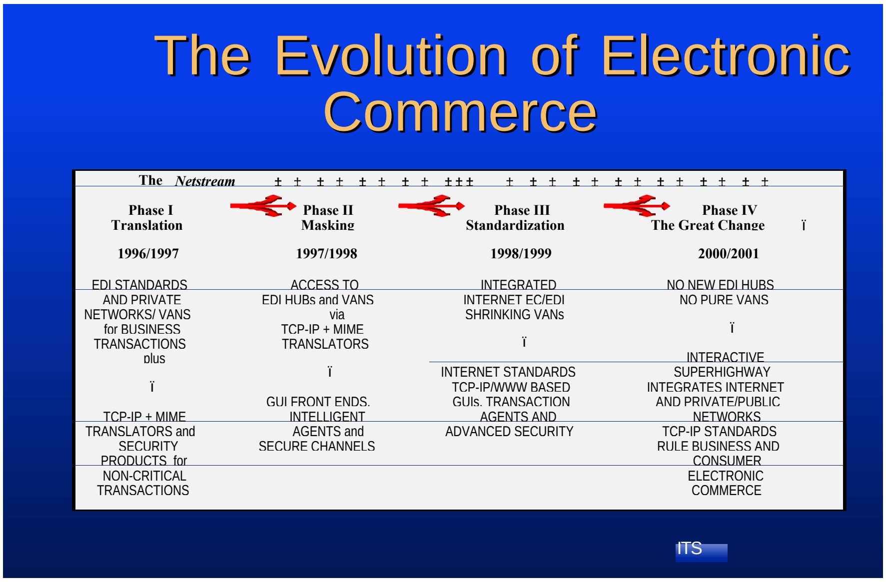

cally with each other and the way they do business electronically with consumers. Today there is a wide divide between EDI as it is carried out between large companies in fixed relationships with pre-defined partners and ECI as it is carried out by "early adopter" consumers who are willing to risk using a credit card to purchase items over the Internet. By the year 2000 this distinction will have been eroded to the point where, even though there will still be systems based on the legacy of EDI back in the bowels of the system, the home or office user will neither perceive, nor need to perceive, the difference.

The Global Market for ECI

The intensity of the focus on consumer-based electronic commerce on the Web is a function not of its current activity (in the "Advertising and Lost Shirts" phase nobody makes any money except through an initial public offering) but of its potential market size. Forecasts usually assume that a transaction is initiated by a PC user who is a consumer at home. However, it is the nature of on-line access to blur the distinction between business use and consumer use. In France, analysis of Minitel usage over the past 10 years suggests that the more appropriate distinction would be functional: a user sitting in her office may access an order-entry system to purchase supplies for her business and to pay her suppliers bills, then go home and order goods from a catalog or a grocery store and pay her personal bills as well. According to a broad definition, this is all ECI.

Travel is the most successful Web-based business according to a survey of more than 1,000 respondents by ActivMedia. The research firm says that 71 percent of travel-related Web sites had sales during the past month. However, total Internet sales represent less than 0.5 percent of travel purchases (*Information Week,* June 1997). Clearly, travel applications cross the business / personal boundary with the majority of revenue coming from business travelers.

Market forecasts vary, as is illustrated below, but the pundits agree on one thing; the market is potentially very large.

- Forrester Research: On-line shopping generated $200 million worth of goods and services in 1994 and will reach $5 billion by 1999.
- Jupiter Communications: Web-based shopping amounted to $132 million in 1994. $4.5 billion will be spent in 2000.
- The Yankee Group: "Interactive consumer commerce" reached $300 million in 1996 and will exceed $3 billion in 1999.
- ActivMedia: Sales generated by the World Wide Web will grow from $436 million in 1995 to $46 billion by 1998. This forecast includes both business and consumer electronic commerce.
- International Data Corp.: Electronic commerce will hit $150 billion in 2000.
- Killen & Associates: E-money transfers on the Internet will reach $600 billion by 2000.
- Input: Business-to-business electronic catalog marketplace will exceed $167 billion in 2000. This is one of the few forecasts that clearly separates the business from the consumer market.
- Nielsen: More than 2.5 million people had purchased products and services over the Web in mid-1995.

• Booz Allen & Hamilton: In 10 years as much as 20 percent of all household expenditures will be funneled through the I-way.

There is optimism about the long term potential of the Internet, or to be more specific, of networks and services that utilize the Internet standards and protocols and their derivatives. However, it may be that the greater short-term impact will be in the business-to-business arena, which will lay the foundation for the global adoption of business-to-consumer services. This insight is based on a long-term view of the cyclical development of interactive services[2] as well as the reality of the global consumer market place.[3]

A few facts as a reality check: Despite the global proliferation of Internet hosts, less than two percent of the world's population has access to the Internet. This may be the same 2 percent that uses a credit card, but it is still only 2 percent. The Internet may be a U.S.-driven phenomenon but commerce is global and it requires global standards and acceptance. International trade and currency issues abound. The Internet is global, whether we like it or not. Companies that offer goods and services for sale over the Internet are offering them to the world. Most U.S. companies are ill prepared to receive, process, fulfill, and support orders from unsolicited customers around the globe.

The incredible growth of access to the Internet in the past two or three years is the result, in large part, of a unique event: the overlay of the World Wide Web on the existing internet infrastructure and the widespread adoption of TCP/IP and related protocols and standards. There will be other such momentous events. However, many believe that the next five years of growth of ECI will be predicated on the enormous acceptance of these standards by businesses for internal and intra-company transactions. This investment will provide the next level of infrastructure, the "commerce layer" on which global consumer interactions, transactions and payments can be securely founded.

References

1. Study of Electronic Commerce, Interactive Telecommunications Services Inc., December 1996.

2. "Open Systems: The Ultimate Distributed System." Hilary B. Thomas, Aregon International, 1984, and "Applying Videotex in Organizations." Hilary Thomas and John Roddy, *Data Processing,* July 1984.

3. ITS Internet Global Monitor, Interactive Telecommunications Services Inc., In process.

Public Key Cryptography: Improved Network Security and Support for E-Commerce

Anthony Wasilewski

Chief Scientist, Digital Video Systems

Scientific-Atlanta, Inc.

Security in broadband networks, including cable and fiber networks, can be enhanced by means of a technology called public key cryptography. Cryptographic methods have been used for many years to secure networks, but many advantages are afforded by using public key approaches instead of secret key approaches. In particular, today's security needs include support of electronic commerce, and public key cryptography is particularly useful in improving this particular application. These technologies do not yet have a wealth of broadly-adopted, standard implementations, but there are some emerging standards that may be highlighted.

The current types of commerce already occurring electronically are fairly broad. Home shopping networks, such as QVC and the Home Shopping Network, where goods are purchased via television advertising and telephone ordering, is about a $3 billion per year industry. On-line shopping, according to the Yankee Group, Boston, could reach $10 billion in consumer transactions by the year 2000. Business-to-business transactions should be considerably higher. This indicates a strong interest by consumers and businesses in shopping through computer networks. Though not yet widely deployed, interactive-TV services could also generate substantial commercial revenues.

Broadband networks include hybrid fiber/coax (HFC), a network in which the trunk line is fiber and very high data rates are possible. Since it is not yet cost-effective in most cases to provide fiber connections all the way to the home-user premises, an HFC network uses traditional coaxial cable to carry the signal over the last several miles or so to the home. Coaxial cable is a broadband medium that supports very high bandwidth. Another broadband network is fiber-to-the-curb (FTTC). In the FTTC approach fiber is taken deeper into the network by deploying termination cabinets near homes, so a node ranging from 8 to 32 homes will be served directly from the fiber line, allowing larger and more symmetrical downstream and upstream bandwidth.

Satellite technology is also evolving. High-powered Ku-band satellites are becoming more popular. The DirecTV service is a one example of the direct broadcast satellite (DBS) approach. This service now has over 2.5 million subscribers and is being called one of the most successful consumer electronics product launches in history in terms of the speed at which users have adopted the service. C band—the traditional big dish —is still popular in rural areas. Finally, terrestrial wireless service offers broadband capabilities: new networks may use multipoint microwave distribution systems (MMDS). This technology has delivered analog cable service for some time, but it is not very widely deployed when compared to HFC networks.

Today, broadband networks are increasing the use of fiber in addition to coax to yield total radio frequency (RF) bandwidths from 750 MHz up to 860 MHz, the latter bandwidth being popular in Europe. Two-way interactive services are being tested, and more of them will be introduced later this decade. Commercially viable digital video and audio compression exist today. The Moving Picture Experts Group (MPEG-2) standard (ISO 13818), which offers broadcast-quality video and audio, is already being used in direct broadcast satellite services. Soon, it will also be used in terrestrial networks.

Access to the Internet via cable modems is becoming very popular as well. In fact, cable operators see this as a valuable incremental revenue stream. Consumers want Internet data faster and less expensively than it can be provided by telephone companies through ISDN or ADSL. It appears that the data products being rolled out have high performance and a very reasonable cost model which should be attractive to both home and business subscribers. Of course, another issue for these operators is connection to the Internet service providers from their head ends and a variety of security and privacy issues.

Routed addressing schemes are becoming popular on broadband networks. In the past, connection-less, "single-wire" addressing to communicate with home terminals was used. In this approach, equipment in the head end merely appends the terminal's address to a message and puts the message on the (usually single) downstream carrier. Then all terminals in the system must filter through all the messages in order to find the one(s) addressed to it. In modern broadband networks, the communications model is becoming more complex.

FIGURE 1
The Digital Security World Is Different

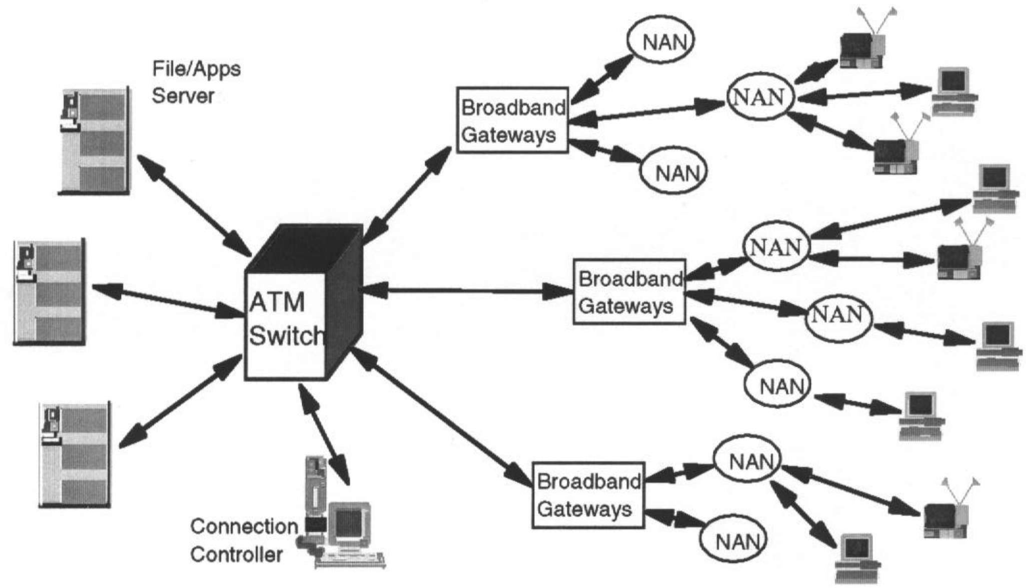

Security demands in digital networks are different from those in previous analog networks. They are being driven by new connection-oriented requirements as the schematic in Figure 1 illustrates. New types of services such as interactive shopping and video-on-demand (VOD) are producing architectures in which an asynchronous transfer mode (ATM) or Synchronous Optical Network (SONET) architecture may be employed as the wide-area transport supporting servers of all types that transmit digital files and video. These can be housed in various regional centers which route information through the network to various types of broadband gateways. That information is routed further through smaller nodes which handle both downstream and upstream information to/from televisions and/or personal computers at the consumer side of the network.

These trends are causing the basic security relationships to change. The current broadcast-only services, such as satellite or traditional cable, use a one-to-many relationship in which the operator of the network establishes an account with each subscriber in the network. This is a preestablished relationship and supports simplified billing on behalf of service providers. This billing is transacted by the network operator. With interactivity, that relationship likely will not continue unchanged. The operator cannot possibly represent every company on the server side and support a scalable and dynamic electronic-commerce environment. The interaction must become many-to-many. Traditional secret key approaches do not support this need as well as public key approaches.

Figure 2 shows what this means in terms of networks. In the center, there is some large-scale transport such as SONET and/or ATM, with various gateways that can bring in digital satellite or off-air signals. This could be, for example, a sports program encoded into digital video format in real-time. Internet services providers could deliver services in such a network as well. After passing through the wide-area network, this transport may be terminated in broadband gateways and processed for further transmission through modulators in the access network. In this example, quadrature amplitude modulation (QAM) is shown, and the signal runs over either HFC or MMDS access networks. Finally, on the subscriber's side, a receiving decoder using MPEG standards is employed. Another relevant standard is that of the Advanced Television Systems Committee (ATSC). Because there are and will be many receivers and senders, support of many-to-many security relationship will become a necessity.

Emerging E-Commerce Applications

Examples of emerging e-commerce services include Internet shopping malls and business-to-business interactions over networks. Electronic cash is also making inroads into the consumer market. Mondex and VISA cash are smart-card forms of electronic cash: *Table 1* lists WWW addresses where more information about these examples is available.

Electronic cash is a stored-value application on a chip card. A chip card is a microprocessor placed in a piece of plastic similar to a credit card with a contact pad installed above the chip so the card can be used in machines. Mondex, originally started by National Westminster Bank, now has 17 major financial organizations involved. Its first trial, started in July 1995 in Swindon, U.K., has been surprisingly successful: hundreds of merchants and over ten thousand users have participated. MasterCard has acquired a 51 percent interest in Mondex which has expanded into many more counties and cities.

Visa cash is another stored-value example. It debuted at the 1996 Summer Olympics Games in Atlanta and was the first

FIGURE 2
Broadband System Architectures

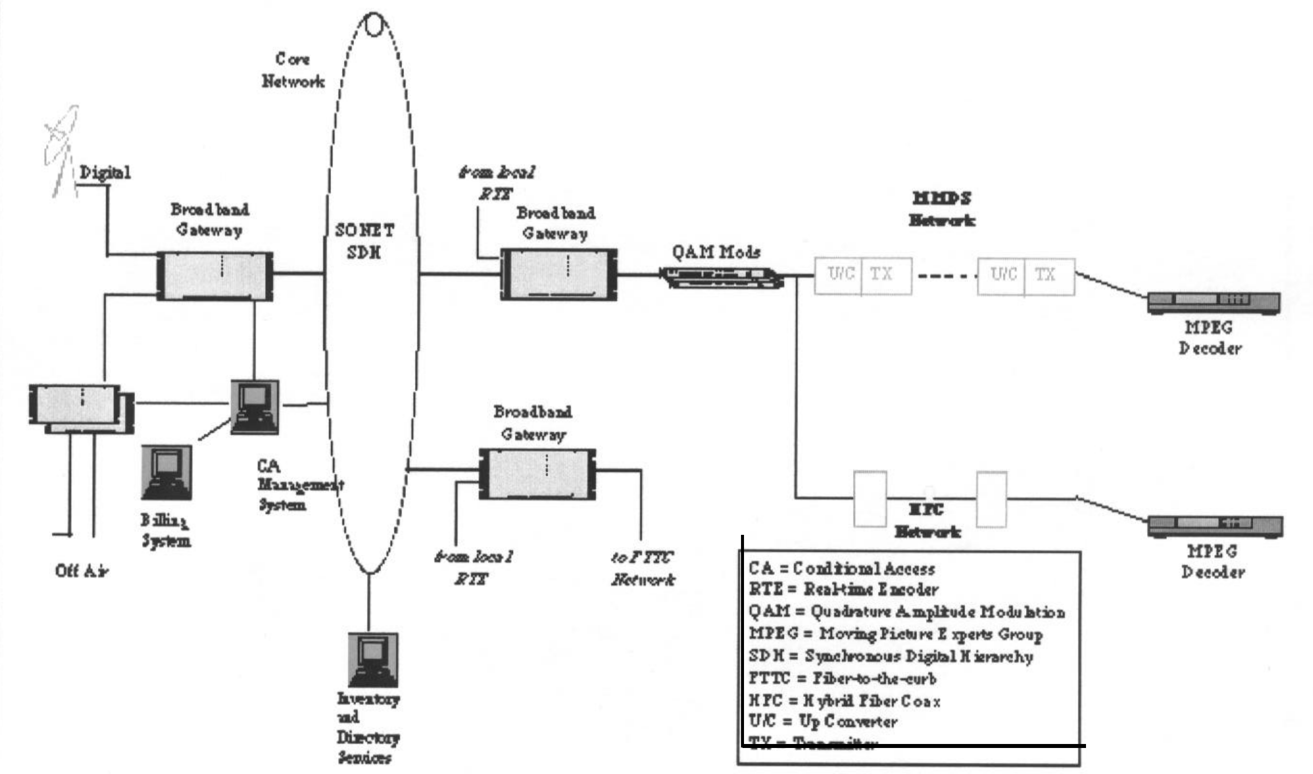

TABLE 1
E-Cash

- MMex
  - http://www.mondex.com
  - stored value on a secure chip-card
  - started by NatWest Bank (U.K.)
  - now owned by 17 major organizations and banks in North America, Asia, Oceania, and Europe
- Visa Cash
  - http://www.visa.com
  - debuted at the 1996 Summer Olympics in Atlanta
  - also, stored value approach
  - $1 million in about 200,000 transactions through end of Games

commercial use of smart cards for electronic cash in North America. By the end of the Olympics, the trial produced more than $1 million in about 200,000 transactions. It worked well where smart card terminals were available. However, these terminals were not always where customers needed them. Nevertheless, it was possible to ride the rapid rail system by just using a smart card instead of fumbling around for cash to buy a token or a ticket. Visa has expanded the plan introducing a total of 18 programs in 13 countries.

Antiquated Security Approaches

It is necessary to know which legacy security methods will not work in advanced broadband networks and understand why. Three examples focus on legacy approaches versus newer security techniques, such as public key cryptography. The first case involves communication over the Internet; the second looks at key management in the cable TV networks; and the third reviews remote access over a network for information systems in a business.

As it has evolved, the Internet has many security weaknesses. Security was not foremost in the minds of the designers of transmission control protocol/Internet protocol (TCP/IP). There is an inherent lack of privacy; all the lower-layer protocols, such as Ethernet, are broadcast, and anyone can peer into them. There is also no basis for authentication because TCP/IP is a connectionless protocol. The user sends packets labeled with the source and destination, and the recipient has to trust that the message really comes from the labeled sender because there is no means of authenticating the sender of the message. In addition, the authentication and integrity of the data itself is in question because other than simple checksums for basic error detection, there is no means to safeguard against malicious tampering. These checksums are so weak that it is very easy to change the data and simply recalculate them.

Key management in cable TV networks is typically also achieved in an unauthenticated manner. Current cable TV systems have security systems and, in some cases, have fairly sophisticated cryptography applied to them. However, because when messaging the home terminals, the same key is required at both the encryption and decryption sites, these secret key databases can be vulnerable to insider attacks. Databases are subject to unauthorized access, and keylists can be stolen during shipping. Using this key information, clones of legal terminals can be made that have the same keys in them. These clone terminals will then respond to messages intended for the legal device, such as the authorization of services, etc. Thus, protection (secrecy) of these home–terminal key databases is of utmost importance.

Employing such a secret terminal key database prevents the distribution of the security function, because distributing the secret keys to many locations makes them potentially more vulnerable. There are more places to protect physically as copies of the key database are replicated. As a result, most current CATV or satellite systems centrally control this function; they house it in one location and maintain close physical security at that place.

To illustrate the third case, imagine a management consultant who works at home who needs access to very sensitive data regarding the company he or she is working for. The client wants to provide the information, but a question arises about ways to do that securely. Exchanging secret keys will enable the transactions, but there are logistical problems with that procedure. Depending on the way secret keys are used, the integrity of the messages themselves may still be questionable. If key management and encryption are not done carefully, it is possible for hackers to enter the system and change the messages. This could result in erroneous data exchange between the consultant and the company or possibly to theft of vital information such as trade secrets.

Public Key Cryptography

Public key cryptography is an approach that can provide many advantages in these areas and fill some of the gaps in security incumbent in older systems. *Figure 3* shows the fundamentals of public key cryptography and contrasts these with secret key cryptography. Public key methods have distinct advantages, however, it should be noted that there are licensing and intellectual property issues that must also be considered. In contrast, many secret key ciphers are in the public domain and may be used royalty-free. There is also a duty to provide digital certificates and a certification authority when using public keys.

Note in *Figure 3* that in a secret key cipher such as the data encryption standard (DES), the same key must be used by the sender and the receiver to encrypt and decrypt the message. The process begins with plain text (i.e., an enencrypted message). The key is then applied—an encryption algorithm such as DES produces cipher text which is sent over a network. At the other end, the same key must be applied by the receiver to recover the plain text again. The implication is that a secure channel is needed to get this key from the sender to the receiver. If the key were transmitted openly, it could be recovered by unauthorized entities through simple network "snooping."

Public/private keys—which are called asymmetric because the encryption and decryption keys are different—are mathematically related to each other. What one key encrypts, the other can decrypt. Anyone with the private key can decipher any message encrypted with the corresponding public key. The RSA algorithm is an example of this approach. Usually, the private key is securely stored so that it may not be easily discovered or altered.

In addition, the public part of the key can be published. It can be well known to all parties and does not have to be kept secret. It can be put in a directory. Knowledge of the public key cannot be used to derive the value of private key. Note that for full security and trust, a certification requirement exists. While the public key can be published, it must be uniquely and securely bound to the identity of its owner for trusted authentication to be possible.

FIGURE 3

The Basics

SECRET KEY

- Encryption and decryption algorithm use SAME KEY
- Thus, key must be transmitted on a secure channel
- Called "symmetric" because the same key is used for all operations
- Called "secret" because the key must be known only to sender and receiver

PUBLIC KEY

- Public Key methods are "asymmetric" because a DIFFERENT key is used for decryption and encryption
- The two keys are mathematically linked
- The Encryption key is "Public", that is, its value may be freely published
- The Decryption (private) key cannot be discovered through knowledge of the Public key

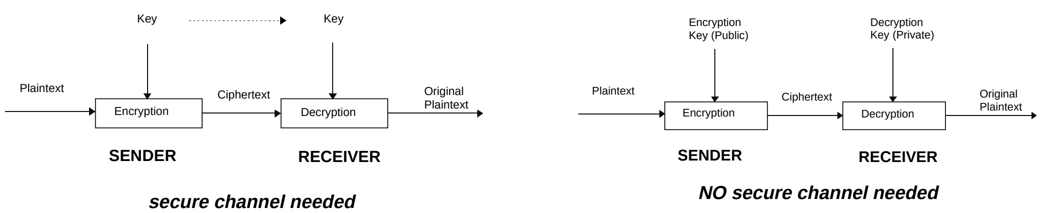

FIGURE 4

How Does Digital Signature Authentication Work?

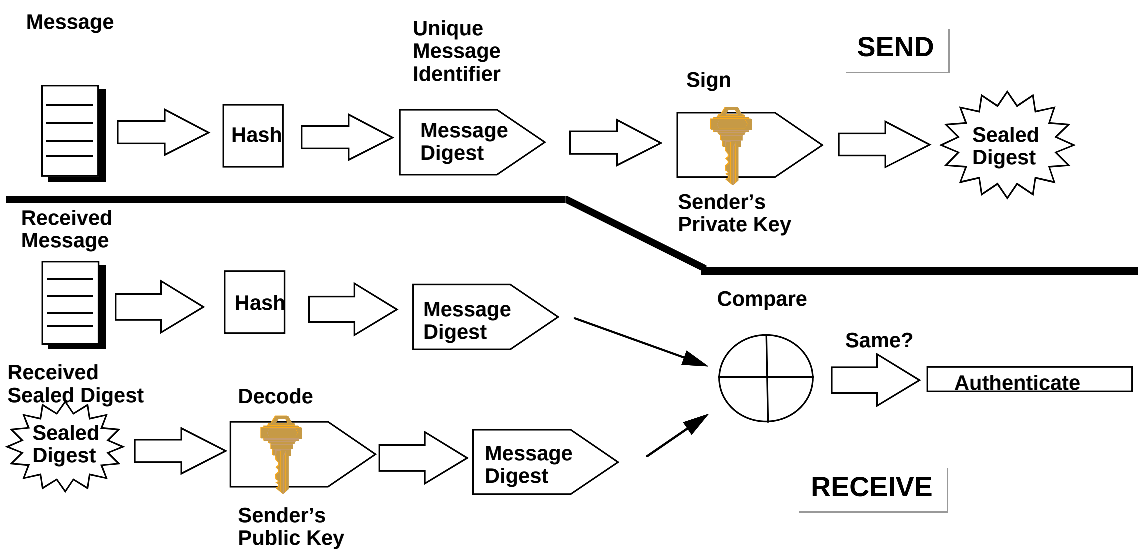

The public key method provides much better support of the multiple service-provider scenario in the sense that no pre-established relationship between each service provider and each user is necessary. Security issues associated with secret key databases are eliminated because the public key can be shared and published openly. Thus, they cannot be stolen because they are already well known. Public-key approaches are the best way to provide a digital signature and authentication of users. This digital signature and authentication allows spontaneous connections between users and service providers or between users in a peer-to-peer mode in which they may not know of each other in advance. If they are registered in a public-key database, they can exchange secure messages and be sure that the other side is truly the entity it purports to be. The sender of the message is authenticated.

Figure 4 illustrates the authentication process. A message—it could be a long message, a file, or maybe just an e-mail—is run through a one-way hash function, which produces a much smaller token (typically 128 or 160 bits) which is called a message digest. This is a unique identifier of that message, and merely knowing the digest does not allow the discovery of the message. In addition, because of the design of the hash function, an adversary cannot even formulate an alternative message which produces the same message digest. This makes an attempt to provide a false message very difficult.

The message digest is then validated by the sender by digitally "signing" it. This is accomplished by encrypting it with the sender's private key. A message encrypted with the private key can then be unraveled with the corresponding public half of the key. Because the sender's private key is used—the one that is not published and that only the sender knows—this provides a way for the sender to put his or her unique digital signature on the message. The signed message digest is also known as a "sealed digest." This sealed digest can be sent along with the message, and the receiving terminal can simply go through the inverse motions to process it: receive the message, decrypt it if necessary, recalculate the message digest, and use the public key of the sender to decode the sealed digest. All the recipient needs to do is compare the transmitted digest with the one it calculated locally based on the received message content. If these two quantities are the same, then the recipient knows two things:

1) the identity of the sender

2) the authenticity of message

Note that the message itself need not be encrypted in order for the digital signature to work.

Secure hash functions require very special design. They should produce output which is very sensitive to every bit of the message. Changing one or two bits in the message should have a very large effect on the message digest itself. Also, it should not be easy to calculate alternate or "colliding" messages which produce the same message digest. An example of a secure, one-way hash function is the Secure Hash Algorithm (SHA-1) which is approved by the U.S. government (FIPS PUB 181-1).

Data integrity and authentication of correspondents are great advantages offered by public key cryptography. However, to trust digital signatures, users must have reliable means of obtaining public keys to use in signature verifications. This is done using digital certificates.

Digital certificates are tamper-proof bindings of a public key and the owner of that key. They usually include a distinguished name related to the owner, perhaps an expiration date, and other data. Trust is established by the Certificate Authority. The Certificate Authority is the entity that applies its digital signature to all entries in the certificate database. Since there are relatively few Certificate Authorities, their public keys can be trusted by publication in open venues, such as newspapers, via software distribution or over the Internet.

Licensing and Intellectual Property Issues

Public key cryptography has several patents associated with its use. Having been invented in 1976, these patents typically have expiration dates ranging from 1997–2000. Licensing of these patents is required for commercial use of public-key technologies.

Usefulness of a Hybrid Approach

As noted, there are many advantages derived from using public key cryptographic methods, but secret key approaches still have important applications. High-speed data, such as Internet data or compressed digital video, still benefit from secret key approaches such as DES because they are much faster. DES, which is a symmetric algorithm, will run perhaps 100 to 1,000 times faster in the same implementation (hardware or software) than a public key algorithm will. The optimum approach is to use both public-key and secret key technologies in the appropriate combination. Such a hybrid approach benefits from the authentication and digital signature capabilities of public key methods in areas such as key exchange and verification of software downloads. At the same time, the speed of secret key algorithms can be exploited to provide confidentiality of large or high-speed data.

Other Issues

When addressing contemporary security requirements public key cryptography should be considered. However, when doing so, several associated issues must also be covered:

1) use / existence of relevant standards
2) operation of Certificate Authorities
3) secure packaging
4) export issues

Standards

There are a few good relevant standards in this technology area but also many that are still under development. The Internet Engineering Task Force (IETF) is actively engaged in this quest. Some of the important methods they are considering include: authentication headers (AH); encapsulated security payload (ESP); and the well-known secure sockets layer (SSL) being promoted by Netscape; secure hypertext transfer protocol (S-HTTP); secure multipurpose Internet mail extensions (S / MIME); and simple key exchange for Internet protocols (SKIP).

The ITU-T X.509 series recommendation already includes a standard for public key certificates. DAVIC, the Digital Audio Visual Council, is about to publish Part 10 of DAVIC 1.2 which will include general security interfaces and tools for multimedia applications. MasterCard and Visa are pioneering e-commerce secure electronic transactions (SET).

The public key cryptography standard (PKCS) is not a true standard. Rather, it is a RSA Laboratories technical note which gives excellent recommendations on how to use public-key techniques. A more formal effort to establish procedures governing the use of public key cryptography is the Institute of Electrical and Electronic Engineers (IEEE) P1363.

Certificate Authorities

Certificate Authorities, which are a very important component of public-key security, are emerging in the commercial sector. Some examples include Verisign, Inc. and the GTE Cybertrust system. It is also possible for a network operator to run its own certificate authority. Public key cryptography requires certification because permitting uncontrolled use of public keys on a network can lead to “spoofing” and masquerade attacks which break down the trust hierarchy.

Physical Packaging

To promote interoperability and retail availability of good security for networks, physical packaging becomes important. For physical security devices, a popular standard is ISO 7816. This is the most universal reference for smart-card technology. It now comprises a series of six parts, covering mechanical, electrical, and protocol interfaces of these hardware tokens.

The Personal Computer Memory Card International Association (PCMCIA) package is also a choice in this area. Costing more than a typical smart card, the PCMCIA package may not work economically for all applications. It can be used for devices such as laptop or personal computers since it fits reasonably well within their cost models. However, for inexpensive terminals, such as Internet access-only terminals or set-top boxes, it may be too expensive unless the consumer can bear some of the cost through retail sales.

Export Issues

If a U.S.-based company is interested in selling products containing security and encryption technologies, it must be aware that the U.S. government regulates export of these technologies. In the United States, cryptographic system exports are controlled by the Department of State and the Department of Commerce guided by the International Traffic in Arms Regulations (ITAR) and advised by the National Security Agency (NSA). While government treatment of this type of product is improving, infamous examples exist in which the government has forced companies to weaken their cryptographic methods, which in some cases, resulted in a swift breakdown of security in those systems.

Conclusion

With good protocols and a combination of secret key and public key approaches, security systems can effectively and safely enable many classes of broadband networks to securely deploy emerging digital services and support e-commerce applications.

Web-Based Interactive Commerce: a Reality

Subbarao V. Wunnava
Professor, Electrical & Computer Engineering Dept.
Florida International University

Michael Woon Choy
Research Associate, Electrical & Computer Engineering Dept.
Florida International University

Conventionally, buyers visit stores and other commercial places to examine items for purchase. Excessive time is spent on traveling and visiting stores, which are not always conveniently located. At Florida International University (FIU), a Web-based interactive commerce solution with operational models is being developed. For example, digitized video files about a consumer item are placed on a Web server with interactive buttons for consumer manipulation. Java applets are used to complete the storefront appearance as well as for security reasons. A consumer from a remote location can access the Web and in near real time can observe the item under consideration. In this article the authors discuss the design methodology and different possible scenarios to make this kind of information available in a real-time, highly secure environment.

Introduction

Electronic Commerce is a way for businesses to advertise and sell products to a wide audience, to place and receive orders, and to undertake joint development, planning, and production. In some areas it is already serving a purpose. In the case of on-line shopping, for instance, where consumers use personal computers to shop, subscribers can get access to information about products whenever they want and place orders with the advertising company. With more than 30 million users today and 90 million projected in the next two years, the Internet is a new way for businesses to establish computer-based resources that can be accessed by consumers as well as business partners around the world.

This paper will outline a case study in which customers can remotely observe an entertainment commodity, like television, through the Web.

Internet Concepts, Common Gateway Interface, and Java Applets

On the Internet, the hypertext transfer protocol (HTTP) is the primary protocol used to distribute information on the Web. The HTTP protocol is stateless: that is, neither client nor server stores information about the state of the other side of an ongoing connection. With a stateless protocol, no information about a transaction is maintained after a transaction is processed. With a stateful protocol, state information is maintained after the transaction is processed. Servers using stateless protocols maintain no information about transactions and processes.

Dynamic behavior is a vital aspect of the Web, usually in response to some event triggered by the user. The result is a custom-built response, tailored around previous interaction with the user. This is made possible by a special mechanism that supports dynamic creation of hypertext mark-up language (HTML) documents. This mechanism is based on the invocation of external applications that run under the auspices of a World Wide Web (WWW) server, called by the browser in the form of an ordinary-looking universal resource locator (URL) on the Web document currently in use. Each of these external applications can be tailored to deliver customized, on-demand HTML documents. This process is shown in *Figure 1.*

There is a key technology underlying the WWW that permits static information to be queried and imported on demand and that also supports custom construction of Web documents “on the fly.” This technology is called the common gateway interface (CGI). Visual basic CGI programs were written to access a Microsoft Access database with updateable entries in order to create dynamic HTML , and therefore a dynamic environment for the users of the system. The application runs under a combination of Apache and Web site Web servers, under Solaris and Windows NT servers respectively.

The database has been so designed as to categorize and classify products that are advertised. All articles of various types have been correlated and cross-referenced in the database. This allows an invitingly dynamic display of related merchandise that should captivate the interest of the shopper, leading to greater usage. In addition, based on a history of browsing/shopping navigation, some artificial intelligence is being considered to further accommodate the user. Neural networks are being considered for classifying the users’ styles of presentation as well as predicting likely choices. For this classification and forecasting, self-organizing Kohonen maps and back propagation algorithms are also under consideration.

Maintaining the State

A shopping cart CGI is normally responsible for keeping the specific state information for each user on the system. Java ap-

FIGURE 1

Dynamic Behavior of HTML Documents Using CGI

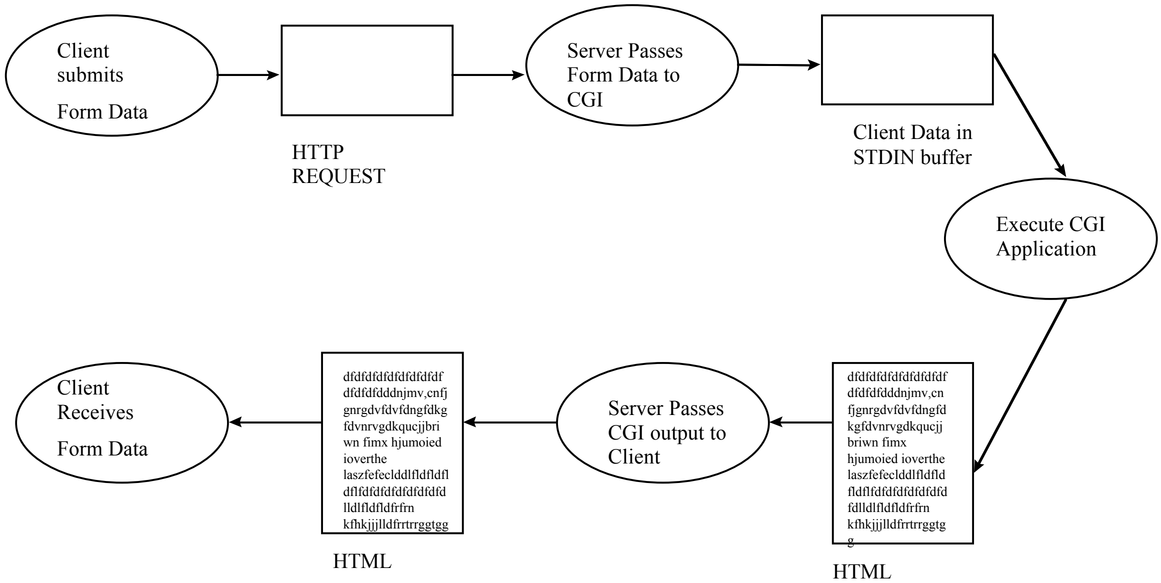

plets were written to provide a better means of keeping state and user information of all shoppers currently on the system. It therefore represents a distributed client-based solution as illustrated in *Figure 2*.

Data Encryption and Security for Transactions

It is essential to implement security measures to ensure that transmissions of highly confidential information such as credit card numbers, etc. are guaranteed. Available security schemes include the following:

S-HTTP

Secure-hypertext transfer protocol is an extension to HTTP that provides security services. It was originally developed by Enterprise Integration Technologies. S-HTTP is designed to provide confidentiality, authenticity, integrity and non-repudiability while supporting multiple key management mechanisms and cryptographic algorithms via option negotiation between the parties involved in each transaction.[1]

S-HTTP uses any of four methods to exchange data-encrypting keys. The possible methods are RSA, out-band, in-band, and Kerberos. If RSA is used, data-encrypting keys are exchanged by the RSA public-key cryptosystem. Out-band refers to an external key agreement, while in-band refers to a key transported in a S-HTTP protected message in another session. Cryptographic algorithms supported by S-HTTP include DES, two-key, and three-key triple data encryption standard (DES), DESX, IDEA, RC-2, and CDMF.[2]

Secure Socket Layer

The secure socket layer (SSL) handshake protocol was developed by Netscape Communications Corporation to provide security and privacy over the Internet. The protocol supports server and client authentication. The SSL protocol is application independent, allowing protocols like HTTP, FTP, and Telnet to be layered transparently on top of it. The SSL protocol is able to negotiate encryption keys as well as authenticate the server before data is exchanged by the higher-level application. The SSL protocol maintains the security and integrity of the transmission channel by using encryption, authentication, and message authentication codes.[1]

A variety of cryptographic algorithms are supported by SSL. During the "handshaking" process, the Rivest, Shamir, and Adleman (RSA) public-key cryptosystem is used. After the exchange of keys, a number of ciphers are used. These include RC2, RC4, IDEA, DES and triple DES. The MD5 message-digest algorithm is also used.[2]

Differences Between SSL and S-HTTP

The main differences between the protocols is the layer at which they operate. SSL operates at the transport layer and mimics the "socket library", while S-HTTP operates at the application layer. Encryption of the transport layer allows SSL to be application-independent, while S-HTTP is limited to the specific software implementing it. The protocols differ with regard to encryption as well, with SSL encrypting the entire communications channel and S-HTTP encrypting each message independently.[1]

Public-Key Cryptography vs. Symmetric Cryptography

RSA and DES are the two most popular public-key and symmetric cryptography schemes, respectively, available today. RSA is about one-thousandth as fast as DES and requires keys about ten times as large. Public-key cryptography and symmetric cryptography are used for different purposes. Symmetric cryptography is used for encrypting data. It is orders of magnitude faster and is not susceptible to chosen-cipher text attacks. Public-key cryptography is best for key management.

FIGURE 2

Distributed Client Based Solution

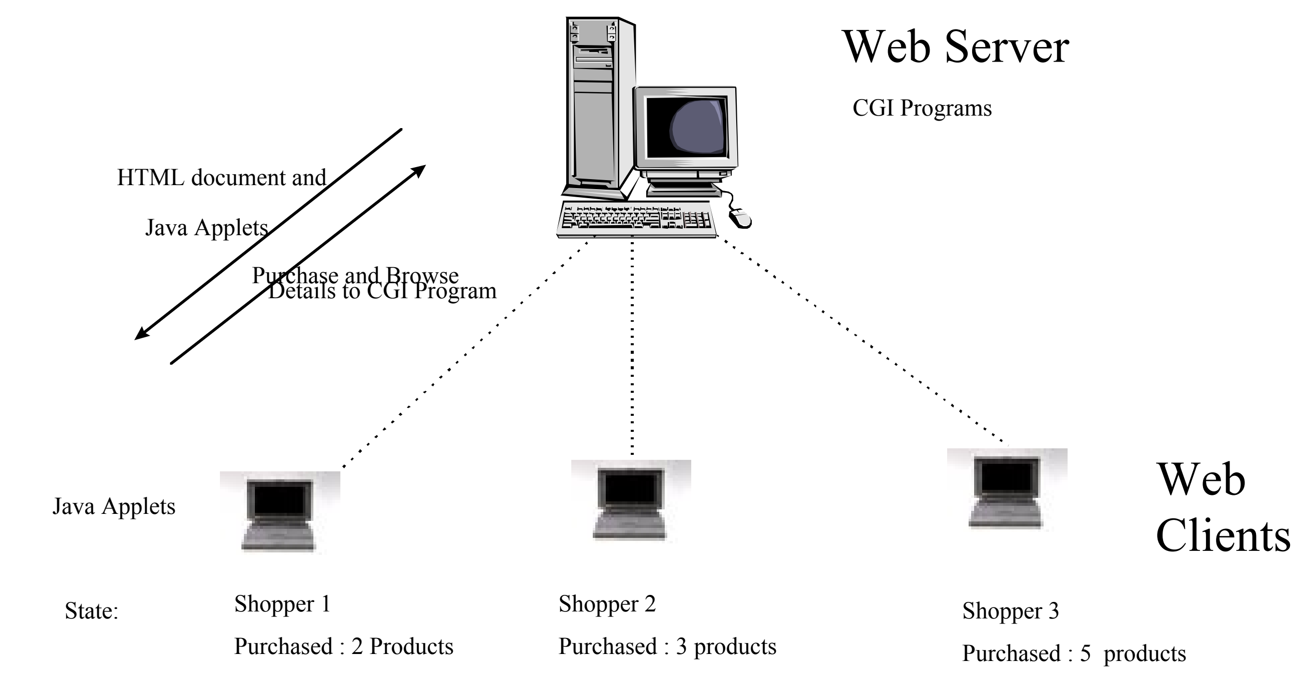

DES is a block cipher, encrypting data in 64-bit blocks. A 64-bit block of plain text goes in one end of the algorithm and a 64-bit block of ciphertext comes out the other end. DES is a symmetric algorithm: The same algorithm and key are used for both encryption and decryption, as shown in *Figure 3*.

Application of Encryption on the Internet for Electronic Commerce

To provide secure transactions for Internet Electronic Commerce, Java applets were written to capture the shopper's demographic information required to complete a purchase.[3] In addition, credit card information and details of the purchase are also captured. This information will be encrypted by the DES algorithm and sent to the server via a CGI script (in addition to the encryption schemes specified earlier). The CGI script on the server end will decrypt the shopper's data and save it in an appropriate database.

The Java applet is issued from the server and contains compiled code. Therefore, the private key shared by the server and the current browser/shopper would not be in clear text. To provide greater security, however, the private key may be encrypted with either the MD5 one-way Hash function or the RSA algorithm. The Java applet would then be used to create a random number, which it uses as the key to the DES encrypting algorithm. This key is then encrypted on the client side and sent to the server to be decrypted, where it can be used by the DES algorithm to decrypt the rest of the data. This process is outlined in *Figure 4*.

Streaming Audio/Video for Product Sampling

Enabling compelling video on the Web required the innovation of streaming video. With streaming video technology, a video is displayed to the end user as it streams over the network in real time. Other than a few seconds latency at the beginning, there is no wait associated with watching a video clip, regardless of length. The user, in general, has the capability to skip around within the video, just as is possible with a music CD.

The Vxtreme product (recently acquired by Microsoft) was used to the do the audio/video streaming for this commerce

FIGURE 3

Symmetric Algorith DES

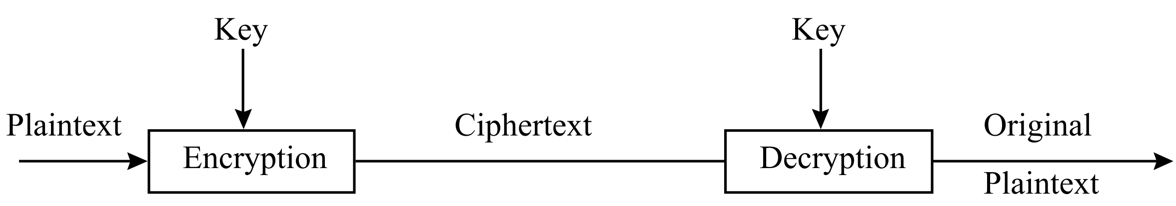

FIGURE 4

Data Encryption Methodologies

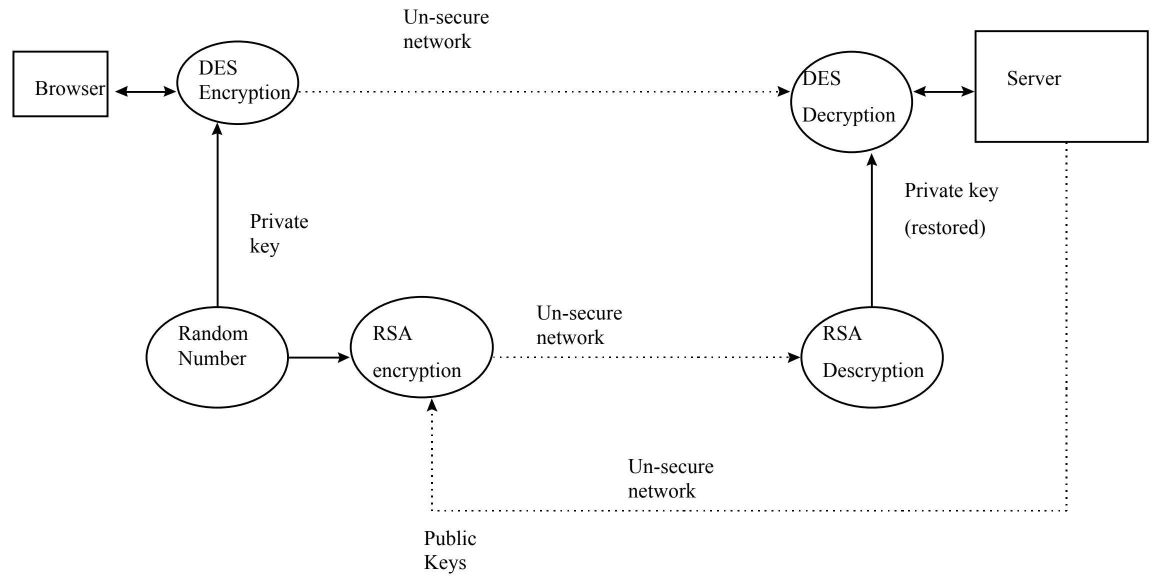

implementation. Its qualities of (a) bandwidth scalability, (b) resolution, frame-rate, frame-quality scalability, (c) fast compression / decompression, (d) ability to cope with network losses, and (e) encoding and decoding latency, proved quite satisfactory for the requirements.

The Vxtreme software uses its own unique compression algorithm and does streaming via real-time protocol (RTP). Unlike transmission control protocol (TCP) or point-to-point protocol (PPP), which can accommodate new field values alone, RTP separates the specification of much of its protocol format into separate files for each type of payload to be transported. RTP uses mixers to allow degraded versions of multimedia multicasts. This is possible because RTP distinguishes between synchronization packets and content packets. Translators are used to funnel multicast streams through firewalls and other constriction points on the network.

Each RTP process periodically multicasts a report giving its name and information like the quality of the audio reception for the participant. Adaptive encode / decode algorithms can then be used to enhance analog audio output at the client computer. For simultaneous audio and video, this is done via separate pairs of ports and data streams. The sequence and time stamp information in the packet headers provides the information needed to synchronize audio and video output to the participant.

Server-Based Video Streaming

Vxtreme uses a media server to serve its media streams, rather than allow the file to be supplied by a standard Web server. One of the main advantages when using a server-based approach is that the server can control the way packets are written to the network. Hence, it is possible to use a specialized protocol built on top of user datagram protocol (UDP) optimized for real-time streaming (RTP) instead of the TCP-based protocol (HTTP) used by the serverless approach. This allows a higher bandwidth to be delivered to the client resulting in better video quality, even when assuming the same network connectivity between server and client and the same level of congestion on the Internet.

By having a specialized video server, the rate at which the data will be consumed, based on headers of the compressed video file, is known, and the bandwidth at which the client is connected to the server (e.g., 28.8 kbps) is often also known. The server sends data to the client only at the required bit-rate, and doesn't drive the network to lose data in the bottlenecked link. Thus, the network throughput is better, resulting in better quality video for the client.[4]

Video Streaming Implementation Procedures

The first step in media streaming was capturing the audio and video to an analog source. This was done using a Sony Hi-8 Camcorder. Depending on the size and location of the products, capture was occasionally done directly to the video and audio capture cards. At other times, it was recorded on tape and played back to these capture cards. The audio / video clips were then stored as audio / video interleaved (AVI) files.

The next step involves using the Vxtreme capture tool to compress the AVI files to VxpressTM format.[4] One audio file is created. Depending on the quality of video and bandwidth requirements, two types of video files may be created, one for local area network (LAN) quality video and one for modem / integrated services digital network (ISDN) quality video. The stream files are then placed on the Web theater server machine. The HTML root page, frame pages, and graphic images are placed on the Web server itself, as shown in *Figure 5.* The stream files are then uploaded into the Web theater server using the Vxtreme author application.

FIGURE 5

Architecture of Web Commerce Solution Using Server-Based Video Streaming

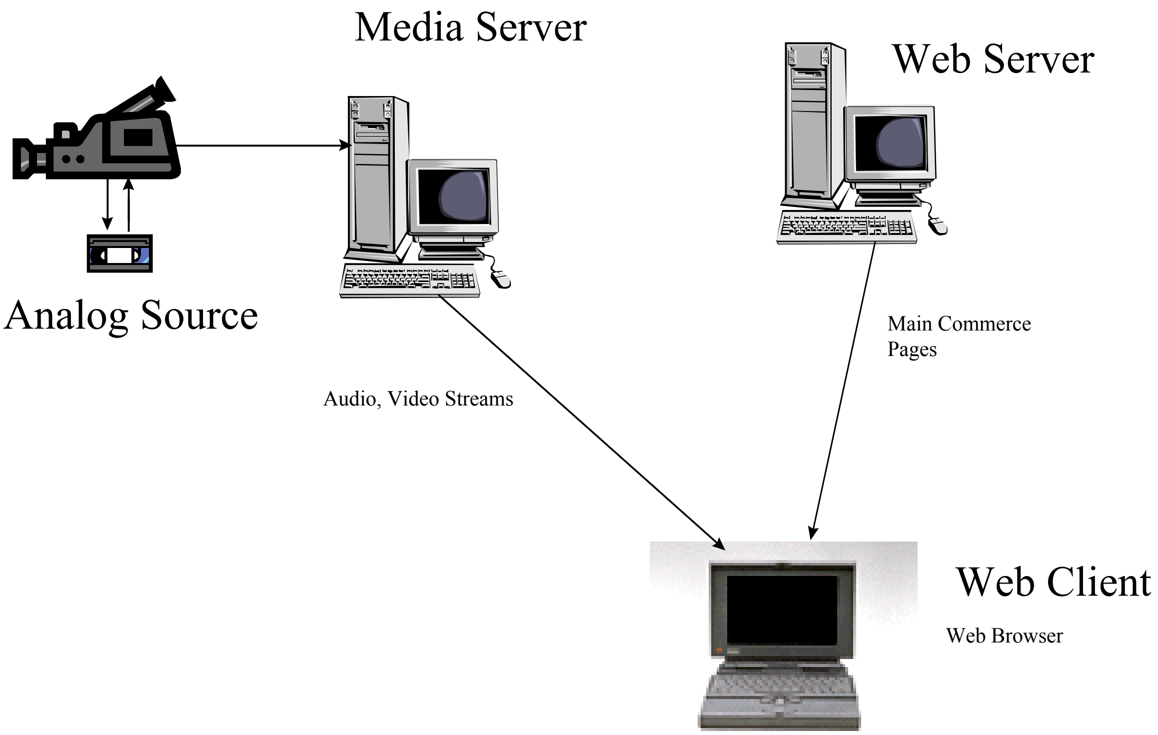

Virtual Reality and 3-D Models

Capture and Conversion of Pictures to 3-D Models

The capture of the pictures for the electronic commerce was done using a Kodak DC-40 digital camera. The PhotoModeler LX 3.0 product was used to convert the pictures to 3-D models, as shown in *Figure 6*. The export of this program gives virtual reality modeling language (VRML) and drawing exchange format (DXF) file outputs.[5]

Use of VRML and DXF Models

VRML is a very recent development, although there were VR systems long before VRML. A number of popular VR software packages exist, each with its own file format. The only common file format that everyone has been able to import is DXF, designed by Autodesk for CAD applications. For VR, DXF has certain drawbacks: it does not specify object hierarchy, material properties, textured maps, or lighting. An essentially new standard VR, available for Internet usage, had to be developed. This language was called VRML.

The Commerce site is equipped with models of the product serving the following functionalities:

a) full rotation of a model in 3-D to examine all facets of the merchandise. Some merchandise uses only photographs of some preset positions for observation, such as front, back, side, and so on. This is made possible using the VRML output of the PhotoModeler program.

b) a fully integratable model with other leading design software, such as AutoCad. The concept here is that the user can download this model and upload it into the design of an existing system like home, kitchen, and so on. This could be particularly useful when purchasing a stove or refrigerator, for example. Being a widely acceptable and easily convertible format, the DXF output of the aforementioned program is used for this.

Implementation of Commerce Application

The various technologies which make this site possible have already been discussed. Diagrams and screen snapshots with browsers like those in the actual implementation are shown in *Figures 7a, 7b,* and *7c.* Navigation through the site is carefully monitored by means of Java applets and CGI programs. This allows the site to have a certain tailored appearance at certain points based upon previous selections.

The product sampling is the critical part of the application, and every effort is made using the aforementioned technologies to make it as realistic as possible. In *Figures 7a, 7b,* and *7c,* a television is being sampled, having been selected from a group of related items, brands, and models. Such features as the tint, color, contrast, etc. can be sampled based on prestored selections of audio/video streams. Each clip lasts about 10 seconds. The display used is 160 x 120 which is comfortable enough to examine the product without losing detail. This has been found to work for the most conservative connections, typically modems.

FIGURE 6

Pictures of Product Taken at Various Angles to Create 3-D Model (Courtesy: PhotoModeler)

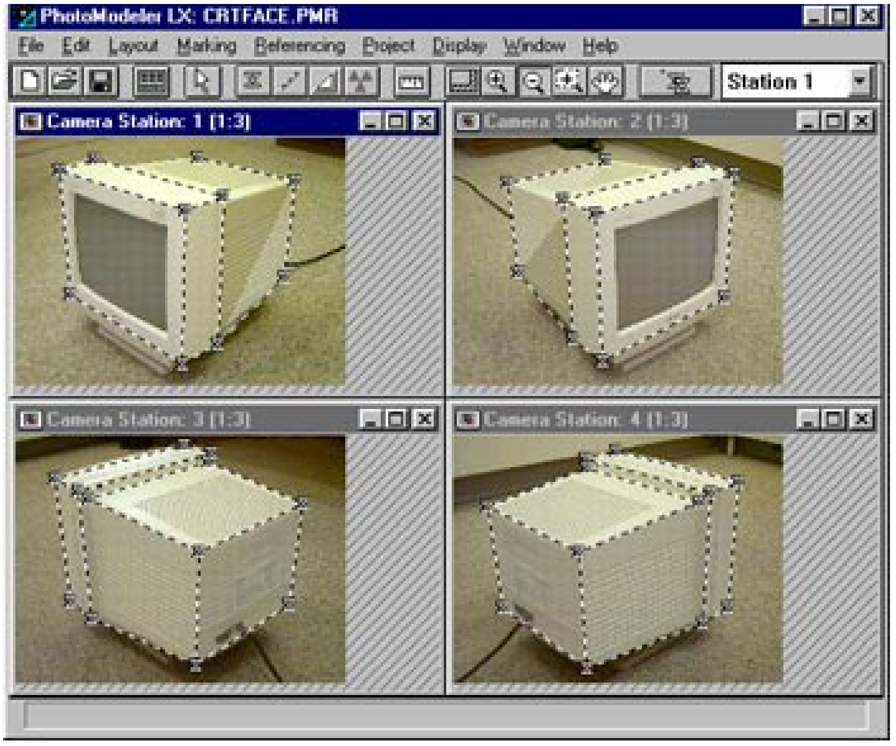

FIGURE 7A

Conceptual Model Used to Manipulate Various Televisions for Examination

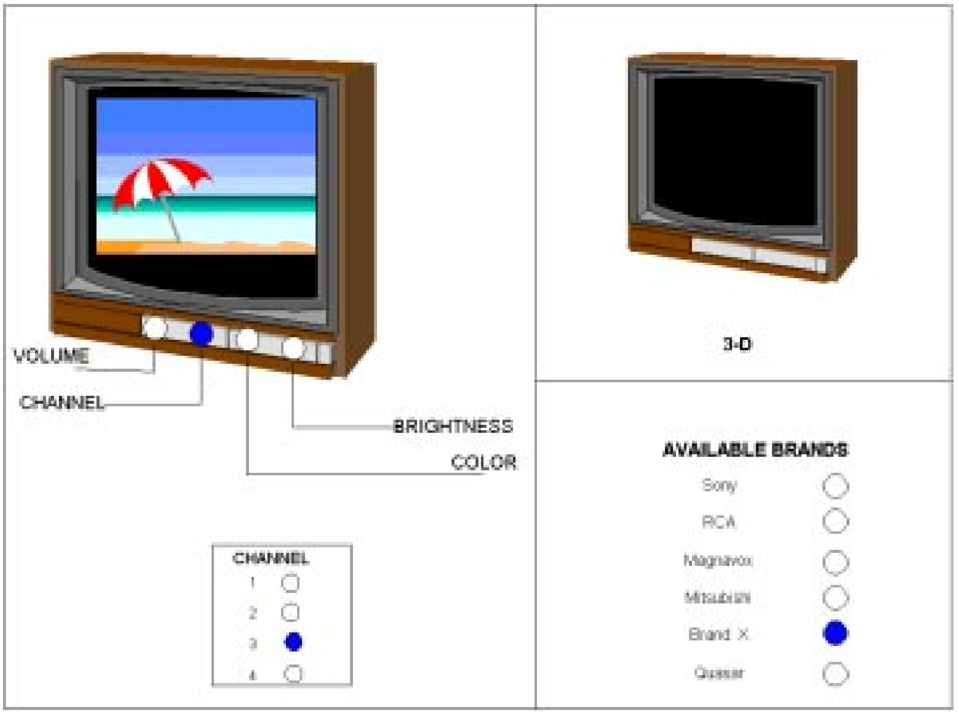

FIGURE 7B

Sample Page in Netscape Used to Manipulate Television for Examination

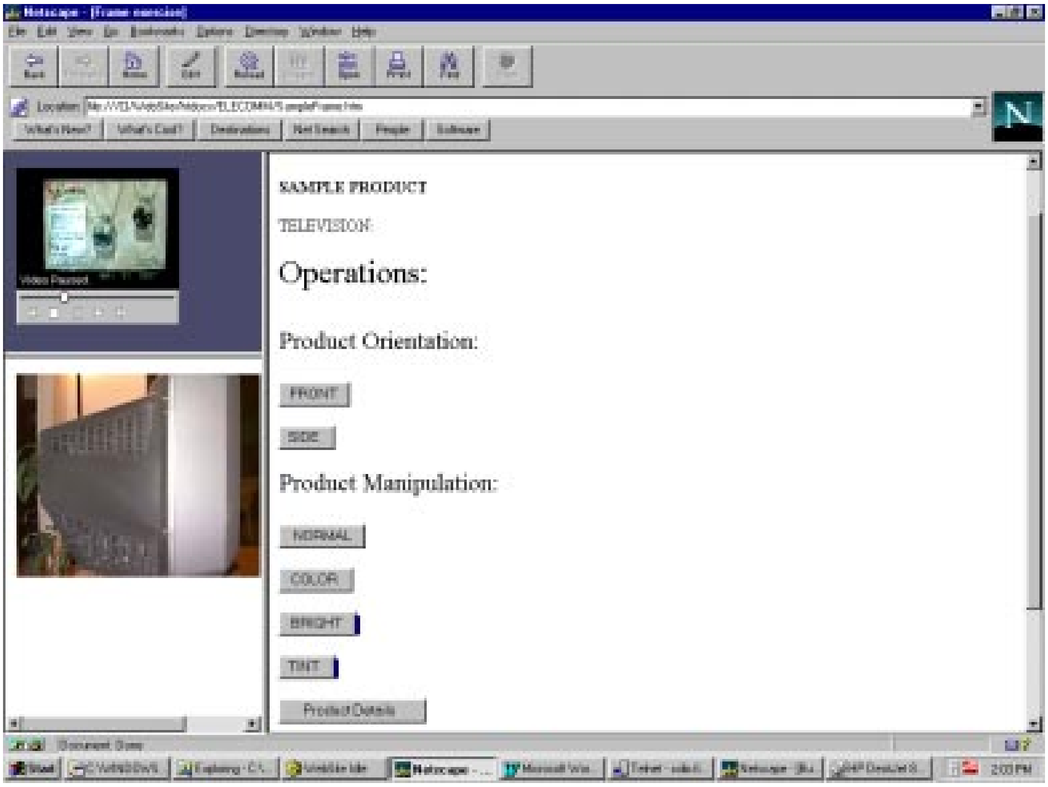

FIGURE 7C

Sample Page in Internet Explorer Used to Manipulate Television for Examination

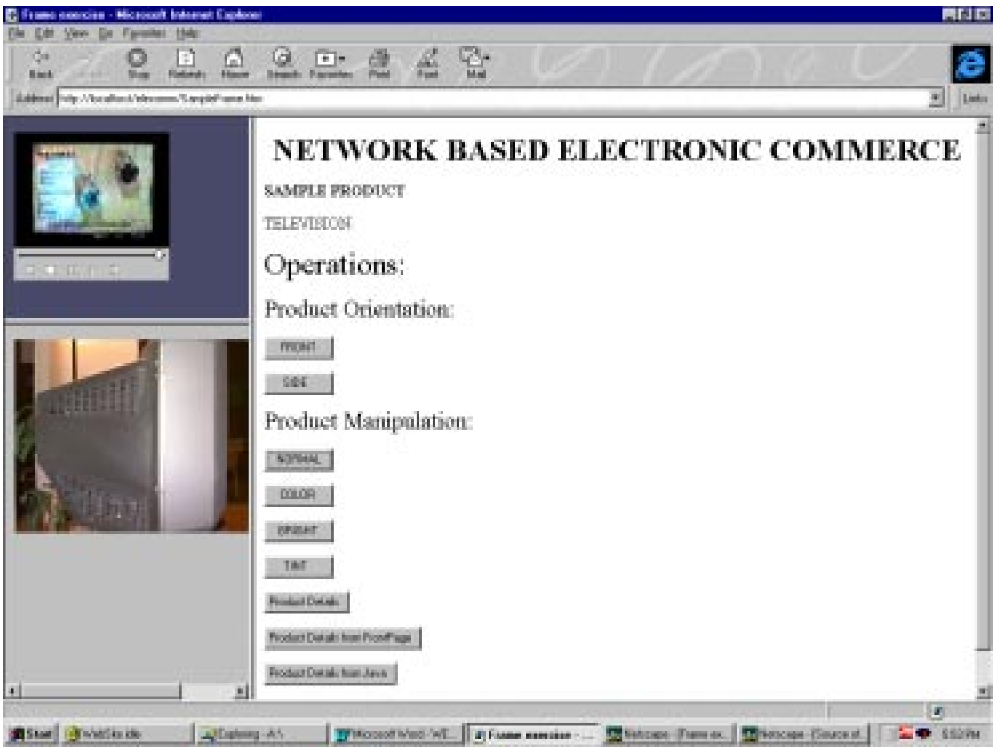

A few products had between four to six pictures taken at different angles. These were used to create 3-D models in VRML (for viewing in Web Browsers) and DXF format (for download into 3-D software). The customer could then rotate the television, for example, through all angles for inspection. It should be noted, however, that a VRML-enabled browser is required for viewing.

Conclusion

The use of CGI and Java applets provides a balance between server-based processing and client processing, respectively. Typically, the CGI programs were used to handle applications where storage and retrieval of data were required. This is not currently possible in the Web browsers available because of certain security restrictions. Java applets were used to perform encryption functions and CGI programs decryption, providing additional security. The tight coupling between the client and media server for media files using Vxtreme resulted in a quality of reproduction to the participant, allowing real-time product evaluation.

Acknowledgments

The authors would like to acknowledge Racal-Datacom for providing the opportunity to learn fundamental cryptographic principles and Dr. Armando Barreto of Florida International University for his concepts regarding the usage of neural networks. Thanks are offered to the members of the Data Communications and Networking Lab at FIU, namely Rizwan Hanif, Madhusudhan Reddy, Manoj Monga, Bruce Moreland, Sridevi Pinapaka, Sanjay Patel, Ajay Sishodia, and Timon Williams for their support , encouragement, and assistance.

References

1. *Answers to Frequently Asked Questions about Today's Cryptography*—RSA Laboratories.
2. *Applied Cryptography*—Bruce Schneier
3. *Java Programming for the Internet*—Michael Thomas, Pratik Patel, Alan Hudson, Donald Ball
4. *http://www.vxtreme.com/*
5. *http://www.photomodeler.com*

Multimedia QoS Service and Management Architecture Using CORBA/Java Platform

Sakir Yucel
Department of Computer and Information Sciences
The University of Delaware

Toshihiko Kusano
NEC Corporation

Tuncay Saydam
Department of Computer and Information Sciences
The University of Delaware

Introduction

Advances in the management of network resources have matured to a level such that it is now possible to construct new value-added logical and virtual services on top of the basic network and management services. Network management is the definition, integration, and coordination of all the hardware, software, and human elements to test, monitor, poll, configure, analyze, evaluate, control, and manage the network resources to meet the real-time, operational performance and quality of service (QoS) requirements. Service management is the creation, access, usage, and management of the value-added services using the logical, virtual service resources as well as the physical network resources and the network management systems. It is important to point out that multimedia service provision and service management, although a vehicle that provides new services is itself a service and therefore must be managed as a logical resource. The separation of a multimedia service, service management, and network resources is crucial in creating open, object-oriented, generic, and transparent service architectures. This paper will address an architecture for the multimedia service and management composed of a generic component. Functionalities of the components are assigned and the interactions among them are developed. Finally, the manner in which the proposed object-oriented service architecture can be realized on a CORBA / Java platform is shown.

TABLE 1

QoS Parameters

General QoS Parameters for Network Applications	Additional QoS Parameters For Multimedia Applications
Throughput	Application Type
Bandwidth	Service Type
End-to-end Delay	Sample Size
Residual Error Rate	Sample Rate
Establishment Delay	Compression Type and Ratio
Establishment Failure Probability	Frame Size
Resilience	Frame Rate
Release Delay	Screen Resolution
Release Failure Probability	Color Depth
Protection	Sound Quality
Priority	Holding Time
Connection Charge	Average Response Time
Data Transfer Charge	Jitter
Utilization	Copyright Charge
Mean Time Between Failure (MTBF)	Synchronization Skew
Mean Time To Repair (MTTR)	Partial Order
QoS Parameters Related to ATM Networks	Partial Reliability
Peak Cell Rate (PCR)	Commitment Level
Peak Duration (PD)	Error Profile
Sustainable Cell Rate (SCR)	Signal To Noise Ratio
Maximum Burst Size (MBS)	
Cell Transfer Delay (CTD)	
Cell Delay Variation (CDV)	
Cell Delay Variation Tolerance (CDVT)	
Cell Loss Ratio (CLR)	
Severely-Errored Cell Block Ratio (SECBR)	
Cell Misinsertion Ratio (CMR)	
Cell Error Ratio (CER)	
Cell Loss Priority (CLP)	

Multimedia QoS Parameters

There are two classes of QoS parameters. The first is the applications' general QoS parameters that contain the common QoS parameters to any application. The second class is the multimedia QoS parameters that also have the parameters more related to the multimedia applications. The first class of QoS parameters is shown in the first column of *Table 1,* while the second class of QoS parameters are listed on the second column of *Table 1*. To satisfy the end-system and networking requirements of multimedia applications and to guarantee the QoS parameters in *Table 1* end-to-end, object-oriented software engineering methodology is employed to develop a QoS management architecture for multimedia services.[1]

Multimedia QoS Service and Management Architecture

The multimedia QoS service and management architecture proposed here provides the multimedia QoS service that is instantiated upon an application's request to access the network resources and which adds values to this service through the management of the service. The multimedia QoS service is composed of basic network services plus the values added by network management services.

Multimedia QoS service creation, access, usage, communication, and management environment may be abstracted through four layers and three planes as shown in *Figure 1.*

FIGURE 1

Multimedia Services in Four Layers and Three Planes

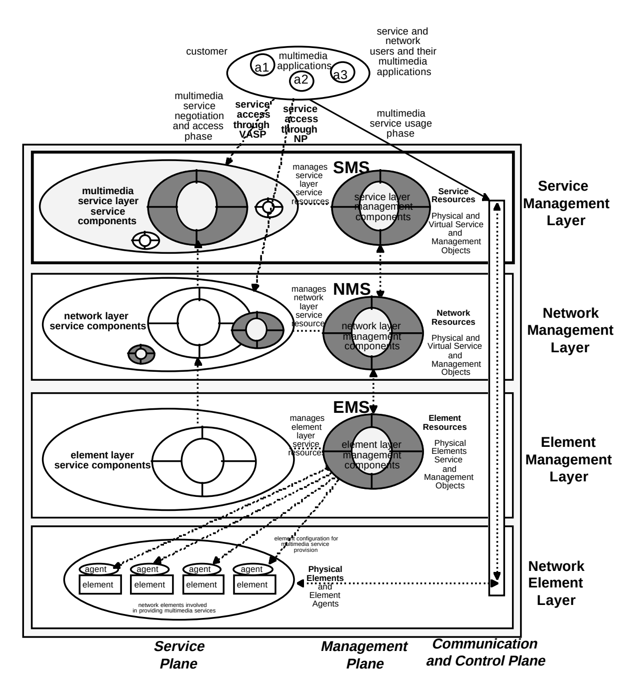

Multimedia applications go through four layers to accomplish peer-to-peer communications: the service management layer, the network management layer, the element management layer, and the network element layer. Each layer has its own logical and/or physical resources which contribute to this communication. The service management layer represents the value-added server provider's (VASP) domain whereas the network management layer is owned by a network provider (NP). VASP is responsible for the end-to-end multimedia service and may span across more than one NP environment. Orthogonal to these layers, there are three distinct planes: the service plane, the management plane, and the communication and control plane. On the service plane are the element, network, and service resources that are needed for the multimedia QoS service. On the management plane are the management systems that manage these resources corresponding to different layers. The element management system (EMS) manages the physical elements that are configured for this multimedia QoS service. The network management system (NMS), on the other hand, manages the network(s) resources that the end-user applications will be using for the multimedia QoS service. Finally, the service management system (SMS) is involved in managing the service resources that actually provide the multimedia QoS service. At the intersection of the planes and the layers are the multimedia QoS components. The communication and control plane has the resources that will be used after the multimedia QoS service is authorized, configured, and accessed.

Interactions between the service management systems and network management systems are essential to instantiation, management, access, and the usage of the QoS service. Hence, the customer environment is connected to both the service management systems and the networks. In fact, a customer may directly connect to network management layer to access the network services provided by an NP without going through a VASP environment. In this case, the customer does not want the value-added services provided by the VASP, but rather is satisfied with the bearer services provided by an NP.

This paper will concentrate mainly on the service management layer. The resources at this layer are primarily composed of multimedia QoS service, service management components, and computational objects, all implemented in software. These resources provide a range of functionalities for multimedia QoS service provisioning and management. Like the network resources, multimedia QoS service and management resources are considered to be distributed. To design an architecture for this service management layer software with the goals of software reusability, inheritance, and flexibility, a doughnut-shaped generic multimedia QoS service component will be developed from which the specific multimedia QoS service component will be developed. The same component will be used to instantiate the multimedia QoS service components at the network management and element management layers. It is also possible to use the same generic component to derive the multimedia QoS management component to manage this generic component.

The Multimedia QoS Service Component

The Generic Multimedia QoS Service Model

The generic multimedia QoS service component is a doughnut shaped Telecommunications Information Networking Architecture Consortium (TINA-C) based component as shown in *Figure 2*. It is composed of four agents and a core. The core contains a component core logic and two management information bases (MIBs): the QoS_Service_MIB, and the QoS_Management_MIB. The QoS_Service_MIB contains all QoS_Service parameters that an application requires with their agreed values. This MIB contains the parameters of *Table 1*,

FIGURE 2

Generic Multimedia QoS Service Component

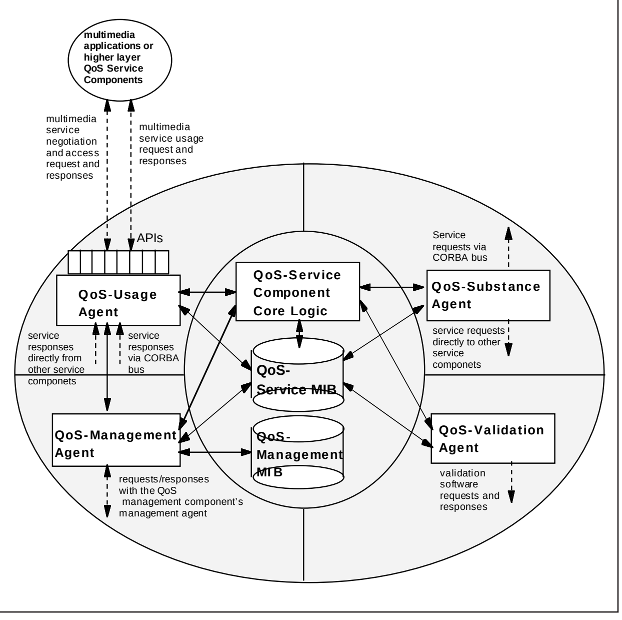

which are mapped onto a component. QoS_Management_MIB contains the QoS parameters that are to be monitored and managed by a specific QoS management component. To implement its core functionalities, that is, what this service component does, the component has generic core logic. The generic component has four agents.

QoS_Usage Agent
This agent has the QoS_Usage interface and QoS_Usage protocol. The usage interface is a set of application programmers interfaces (APIs) that the clients of this component can invoke. Applications as well as other components can access and request the services of this component through the QoS_Usage agent to obtain what this service component offers. The QoS_Usage agent provides a service interface to other components, gets the QoS_Service_Request_Response parameters from components or from the multimedia applications and stores those parameters onto the QoS_Service MIB. Also, this agent gets the service responses from the lower-layer QoS_Usage agents or from the other service components through its interface.

QoS_Substance Agent
This agent has the QoS_Substance interface and provides the component to request the services of other components to help enhance and implement its core functionalities. The QoS_Substance agent is responsible for requests. It should be noted that these are one-way requests and the responses will arrive through the service usage interface. In a sense, with this agent, a generic component, in a client mode, is asking other servers what this service component needs.

QoS_Management Agent
This agent has the QoS_Management interface and QoS_Management protocols. The component service managed by this agent relates to how this service component is managed. All management related information flows across the interface of this agent. The QoS_Management agent manages the QoS_Service_Component, interacts with the QoS management component's management agent through response/request primitives, and accesses and manages both the QoS_Service and QoS_Management MIBs. It monitors and polices QoS_Service_Request_Response parameters in both the QoS_Service and the QoS_Management MIBs. It manages connection requests from other components and controls the admission. It is also the responsibility of QoS_Management agent to estimate resource capacity and generate QoS alarms for QoS degradation on the monitored QoS_Service_Request_Response parameters.

QoS_Validation Agent
This agent with its QoS_Validation interface is intended for a real-time validation (when needed) of the component.

The QoS service component core logic verifies the QoS_Service_Request_Response parameters received by the QoS_Usage agent, optimizes them, and maps them onto the component. It also identifies the QoS parameters that are to be mapped to the lower layer component with their acceptable values.[2]

Armed with these powerful interfaces and with its core functionalities, the component is now ready to offer its own services, form clusters with other components to provide collection of services, or even form a part of many components of a more complex service environment. The same component model was applied to some other core service components: the connection service component, which provides the connection for the multimedia application; the security service component, which configures and activates the security service and mechanisms according to the security requirements of the application, and the bandwidth service component, which provides dynamic bandwidth service. A component using the same model was also designed, called the service interface component, which interacts with the application to create, configure, and activate the multimedia service.

Functionalities of QoS Service and Management Components

This section will assign the functionalities of each QoS service and QoS management component at each layer of *Figure 1* to satisfy the operating system (OS) and networking requirements of multimedia applications. At the same time, it will show how to map the QoS parameters taken from the application onto the specific layers and state which parameters should be controlled by which layer of the architecture.

The Service_Layer_QoS_Service_Component does the following:

- gets QoS_Service_Request_Response parameters from the multimedia application
- maps the related parameters onto its QoS_Service_MIB including *(from Table 1)* screen resolution, color depth, frame size, frame rate, sample size, sample rate, signal to noise ratio, audio quality, holding time, application type, end node synchronization skew, compression type and rate, commitment level, average response time, throughput, bandwidth, end-to-end delay, jitter, partial reliability, security, priority, billing charges, connection establishment delay, connection establishment failure probability, connection release delay, connection release failure probability, and so on
- sends QoS_Service_Request_Response parameters to the service management system (SMS) to start negotiating
- sends QoS_Service_Request_Response parameters onto Network_Layer_QoS_Service component for multimedia service access, negotiation and usage
- indicates the level of commitment (deterministic, probabilistic, or best-effort) to user application depending on the negotiation of QoS_Service_Request_Response parameters
- informs user application(s) in case of any QoS degradation
- indicates service termination to user application at the end of multimedia service usage

- indicates the level of providing for the QoS parameters
- is responsible to provide synchronization for each end-system participating in the multimedia service during the multimedia service usage so that all end systems get the same data within the same interval
- may interact with end-system OSs for buffer allocation, thread scheduling, central processing unit (CPU) and disk I/O scheduling in real time, jitter correction, media synchronization, and flow rate regulation

The Service_Layer_QoS_Mgmt_Component has the following functionalities:

- to control the admission, the connection, the service configuration, and the service reconfiguration during the service access, the service usage and the service reconfiguration phases
- to monitor, control and modify some of the QoS_Service_Request_Response parameters that are mapped onto service QoS component such as holding time, service type, commitment level, average response time, throughput, bandwidth, end-to-end delay, jitter, partial reliability, security, priority, billing charges, connection establishment delay, connection establishment failure probability, release delay, release failure probability, end node sync skew, and so on
- to start renegotiating if needed or requested by the application, possibly with a different set of parameters
- to police application flow rate
- to send alarm to Service_Layer_QoS_Service component in case of QoS degradation
- to estimate resource capacity

The Network_Layer_QoS_Service_Component has the following functionalities:

- to get QoS_Service_Request_Response parameters from the Service_Layer_QoS_Service component
- to map the related parameters onto its QoS_Service_MIB including *(from Table 1)* service type, sync skew, partial reliability, error profile, billing charge (data transfer charge), commitment level, peak cell rate, peak duration, end-to-end delay, maximum burst size (MBS), cell loss ratio (CLR), cell insertion ratio (CIR), cell delay variation (CDV), bit error rate (BER), utilization, mean time between failure, mean time to repair, bandwidth, packet error rate, priority, and so on
- to send the parameters onto the Element Layer_QoS_Service component

The Network_Layer_QoS_Mgmt_Component does the following:

- admission and connection control
- monitoring and controlling the parameters mapped onto the Network_Layer_QoS_Service_Component
- monitoring and controlling multicasting and network congestion level
- monitoring and controlling resource reservation and priority queuing
- resource allocation and de-allocation

The Element_Layer_QoS_Service_Component has the following functionalities:

- to get QoS_Service_Request_Response parameters from the Network_Layer_QoS_Service component
- to map the related network element objects to the component's MIBs and send those objects to the Element_Layer_QoS_Mgmt_Component for management

The Element_Layer_QoS_Mgmt_Component has the following functionalities:

- to configure the network elements
- to poll, monitor, and manage the network element objects that are passed by the Element_Layer_QoS_Mgmt_Component to satisfy the network management layer QoS _ Service _ Request _ Response parameters
- to manage physical resources contributing to the service requests
- to control call admission and manage connections
- to collect the accounting information and pass to Network_Layer_QoS_Mgmt_Component

The API functions corresponding to the functionalities listed above for each component have now been developed. The components do have the same parameters stored in their MIBs, but their views on those parameters are different. For example, the Service_Layer_QoS_Service_Component has the end-to-end view and the value of establishment delay represents the end-to-end value whereas the Network _ Layer _ QoS_Service_Component has the view of its NP network and the establishment delay there indicates the delay within the NP network only.[2]

Interactions Between Service and Management Components

Figure 3 shows how this component concept works by focusing the component interactions between service and management planes. Common object request broker architecture (CORBA) service buses are introduced for communications between multimedia applications and service providers such as VASPs or NPs, and also for the communications between the service and the network layer. Also introduced is WWW-

FIGURE 3

Interactions of Service and Management QoS Components Using CORBA Services

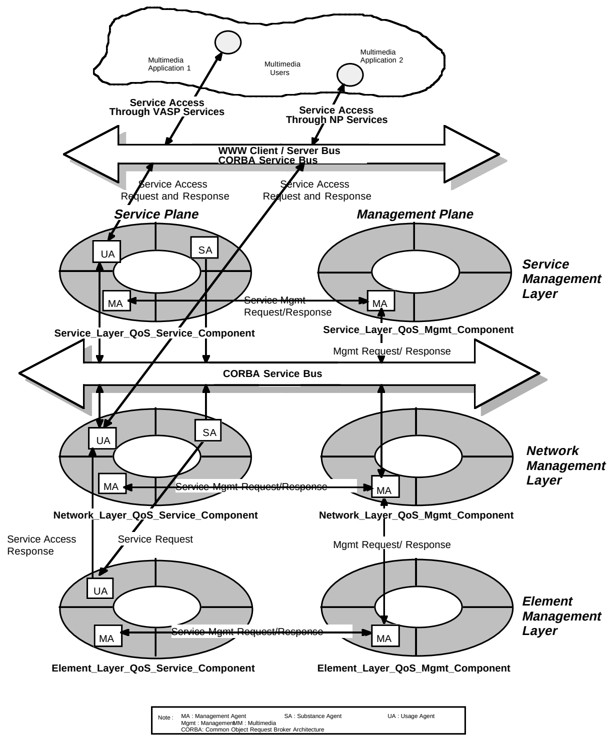

based client-server service on the CORBA service bus for communication between multimedia application and service provider for ease of client / server services. Directions of interactions are categorized into two types:

- inter-layer communications between layers on the same plane
    - upper layer to lower layer is the direction for requesting service
    - lower layer to upper layer is the direction for informing response
- intra-layer communications between peer level components.

The first type is used for service request / response in the service plane. This interaction will appear when the QoS service component requests a service that is offered by a lower layer component. This will utilize the usage protocol by invoking the service APIs of the requested component. Similar interactions occur between management components on the management plane but these kinds of interactions use the management protocol. In both of these cases, the component at the higher layer is the client of the component at the lower layer. The CORBA service bus between service and network layer offers client / server interactions for the components. The second type occurs between a QoS service component and its peer QoS management component. A QoS service component sends the QoS parameters to the QoS management component at the same layer.

Multimedia applications send service access requests to the Service_Layer_QoS_Service Component or to the Network_Layer_QoS_Service Component via a Web-based client / server service supported by the CORBA service bus. By using this bus, multimedia applications do not have to worry about where the server is located or which server has the corresponding objects to realize the requesting services. The QoS_Service_Component registers service requests from applications or other components and sends them as service management requests to the QoS_Mgmt_Component in the same layer to seek for service management help, acceptance, and satisfaction of management requirements. This request is intra-layer communication and management parameters are encapsulated in the message.

The QoS_Mgmt_Component evaluates requests with respect to its management perspectives and it sends a response. The Service_Layer_QoS_Service_Component, then, sends a service access request to the Network_Layer_QoS_Service_Component with network QoS parameters. Between these layers, a CORBA service bus exists and the Service_Layer_QoS_Service_Component, being the client, just asks the CORBA service bus to carry the request to the Network_Layer QoS_Service_Component. Since it is considered that the client in the service layer and the server in the network layer are distributed, it is convenient to use CORBA between these layers. Similar interactions exist between the network and the element layers. Since the element layer is managing the elements, the Element_Layer_QoS_Mgmt_Component takes action to configure the elements by communicating with their agents in accordance with the service agreement. After the Element_Layer_QoS_Mgmt_Component receives the response from agents, it sends management response to the Element_Layer QoS_Service_Component as well as reporting a change of network configuration by management response to the Network_Layer_QoS_Mgmt_Component. In both planes, similar responses will propagate up to the service management layer where the application is notified about the service configuration. Now applications realize their service requests are settled and start using the services.

The main theoretical results with this architecture can be summarized as follows. The architecture is based on the open standards telecommunication management network (TMN), TINA-C, and CORBA. Hence, it is open, scalable, and expendable to various services. It is composed of modular components; even though the component functionalities may differ, the structure of all components remains the same. There is a clear separation of the network element layer, the network layer, and the service layer, as well as between the service and management operations. This separation ensures modularity, transparency, and independence in implementation of services from that of managers.

Since component model structure is consistent throughout the architecture, software reusability is promoted during development. Since the component agents, core logic, and MIBs are thought of as objects, object-oriented analysis, design, and implementation comes naturally. In fact, the service and management components discussed here have been designed and implemented inheriting from the same component object model. Component interfaces are consistently developed to facilitate the component interactions. Even though it can be a design issue, the architecture is so developed as to yield a minimum number of component interactions, thereby minimizing the inter-component communication overhead.

The architecture is conceived to be distributed. The components at one layer communicate with the components on another layer across an open CORBA platform. The architecture enables the applications to access it either through a service provider at the service management layer or through a network provider at the network management layer, in each case using the same APIs over the CORBA interface. It also provides flexible specification and negotiation of service parameters through the standard APIs.

There are some important proposed QoS architectures. The architecture herein described can be compared to several others.[3,4,5] A discussion and comparison of some of the existing QoS architectures such as the XRM, the QoS-A, the TINA QoS Framework, the MASI, and so on is also available.[5] A basic and unique difference between the architecture discussed here and the others is that in this architecture, the QoS service and other core services are controlled and managed by configuring, setting, monitoring, and controlling the MIB objects that reside in the component. At the lowest layer, the objects are mapped to actual asynchronous transfer mode (ATM) MIB objects at the ATM switches. The proper management of the actual ATM MIB objects at the element layer gives rise to the control of the QoS and other service parameters at the higher layers, which, in turn, results in the end-to-end QoS Service guarantee. This architecture thus realizes itself through pure management of the MIB objects at different layers. This requires the proper cooperation of the components as well as the mapping of the service parameters and the MIB objects.[2]

Customer and VASP Interactions on CORBA/Java Platform

This section will show how a user application accesses the services of VASP over the CORBA bus using a Java programming language interface. In the architecture discussed here, the VASP runs an Internet inter-object request broker (ORB) protocol (IIOP) enabled server and the service interface component and core service components are registered CORBA objects providing the interface for the customers to access, create, and use the services. CORBA is a client / server distributed object infrastructure that has become the industry standard as a distributed object platform. A CORBA bus allows the objects to locate and use the services of other objects transparently with regard to location, implementation language, OS and hardware platform.[6] It also enables objects to discover each other during runtime. CORBA defines the interfaces using a language called interface definition language (IDL). The VASP running an ORB server has an open architecture and open interface to any objects willing to access it. What VASPs need now is a graphical user interface (GUI) that allows the customers to select the services they require and to customize both their environments and the services they require. This is what Java is most suitable for, as it generates mobile code that can move around through a special kind of application called an applet and provides a robust, platform independent GUI package.[7] In the implementation outlined here, Java programming language is used not only for programming the user interface but also for implementing the service components discussed above. Integration of Java and CORBA brings together the client-server paradigm, distributed object infrastructure, and portable and mobile code structures which leads the current Web to the Object Web.[8] *Figure 4* shows the interactions between the customer and the VASP server as well as the protocol stacks for both customer and VASP environments. The customer connects to the http server of the VASP using a Web browser and downloads the html page that contains the service creation applet using the http protocol over TCP / IP. This applet is actually the client copy of the service interface component. This Java applet is loaded into the Web browser's Java virtual machine and, having gone through a verification process, starts running in a safe environment.[9] The applet initializes its ORB client environment. Indeed, the applet has the IDL-generated setups that enable connection to the particular server that is the server copy of the service interface component object in the VASP environment. From now on, the service creation applet and the service interface component object communicate by using the usage protocol mentioned above. The binding protocol for this communication is IIOP, which is the protocol to run CORBA over the Internet. The service creation applet has the GUI components that allow the customer to call the service APIs of the service interface component. The customer first selects a specific service such as video-on-demand and fills in the necessary information about that service. Then the customer may want to add some optional core services such as QoS service, security service, and bandwidth service to add more value to his or her specific service. *Figure 5* shows a screen capture for the QoS parameters specification to add the optional QoS service.

When the customer completes configuring his service, the applet calls the CreateService API primitive of the service interface component to obtain the service contract and to guarantee the network configuration all the way down to network elements. The service interface component interacts with the core service components to configure the service objects in the VASP environment and to configure the NP networks that will participate in this particular service. In the implementation outlined here, VisiBroker for Java is used for

FIGURE 4

The Protocol Stacks at Service Client (Consumer) and Service Provider Environments

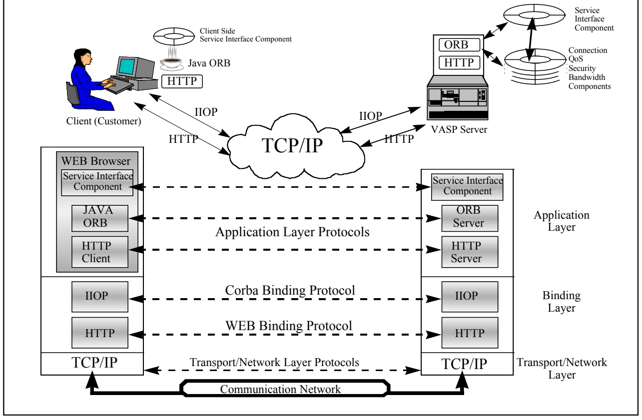

FIGURE 5

QoS Parameters Specification and Negotiation Through Service Interface Component

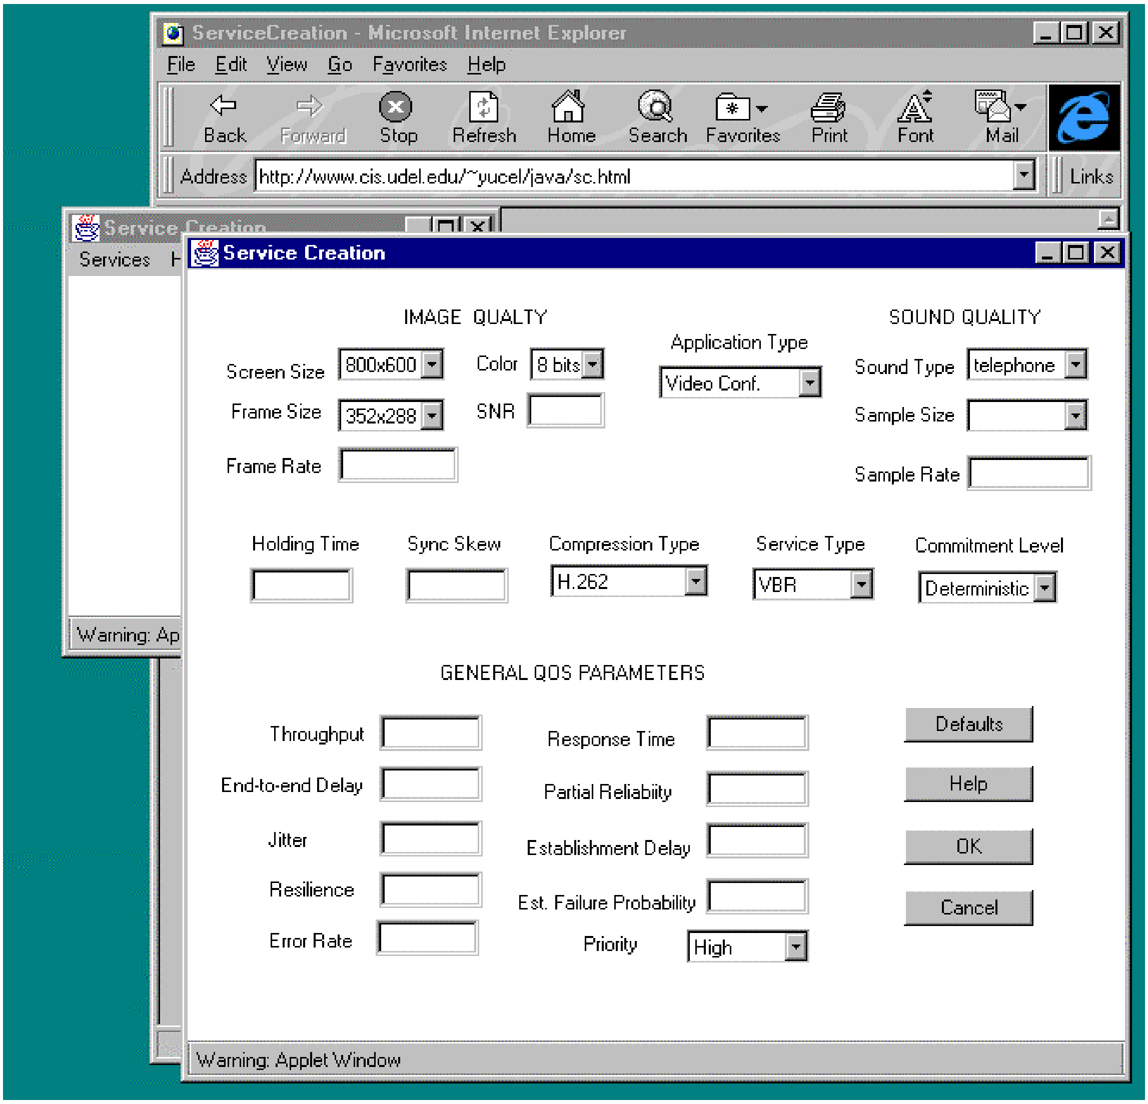

ORB server and IDL/Java integration since it provides both client- and server- side Java support, dynamic method invocation, and Java mobility over IIOP.[10]

In developing the prototype discussed in this paper, an object-oriented software engineering methodology, Fusion,[11] is employed for the specification and design of the architecture, as it facilitates deriving the individual components from the same abstract component model nicely as well as reducing the complexity of the architecture. However, since the distributed software components are supposed to run asynchronously in real time with different timing and QoS constraints, part of the research supporting this paper concentrates on adapting Fusion for real-time and asynchronous systems.

Equally as important as specifying and designing the individual components is to specify formally and implement precisely the inter-component protocol. For this purpose, the authors propose Estelle-like features for the life-cycle expressions in Fusion and for the specification of the inter-component interaction scenarios and protocols.

Java programming language has been found to be very natural for mapping the design object into the program classes, especially in terms of implementing the asynchronous and spontaneous events, threads, exceptions, and timing relations. A further need is to develop methodology in mapping from the design descriptions into Java classes to ease the implementation and the maintenance phases.

Conclusions

Network users are now more interested in services offered to their applications than the underlying communication networks across which the applications communicate. Offered by service providers, multimedia services are critical resources that need to be created, maintained, and managed. In order to be seamlessly integrated, services and the networks must follow the same abstract architectural view of management. The architecture provided here is a coherent and integrated framework for provisioning and managing multimedia QoS services. It is compatible with TMN and TINA-C frameworks. It is integrated, flexible, and transparent to the applications and application programmers. It separates the service access and usage from the management of the service. It also shows a concrete adaptability with distributed management environment by using Java and CORBA services as an implementation method.

This architecture uses object-oriented software engineering techniques to specify and design the components of the architecture to build up the running prototype. Indeed, a generic multimedia QoS service component has been designed, and through inheritance, all of the specific layer multimedia QoS service components have been derived. That same component structure can be used to model and design any component in the architecture, from a single component to a complex service made of imbedded components. The open architecture constructed as well as the generic QoS component models discussed promote distribution and technology independence as well as the separation of logical service environment from physical machine and network environments.

This is a viable, consistent, coherent, and smart way of building open service architectures. In addition, by using Java and CORBA service as an interface implementation, this architecture truly maintains the openness and the interoperability to the customers of multimedia services.

References

1. Yucel, S., T. Kusano, and T. Saydam, “The Generic Service Management Architecture,” *Proceedings*, Vol.1, ATIRP Conference, January 1997.
2. Kusano, T., T. Saydam, and S Yucel, “Mapping QoS Parameters for Multimedia Services Using ATM MIB Objects,” submitted to ICC ’98 Atlanta, Georgia.
3. Campbell, A. et al, "A Quality of Service Architecture," *ACM Computer Communications Review*, 1994.
4. Nahrstadt, K. and J. M. Smith, “The QoS Broker,” *IEEE Multimedia*, Spring 1995.
5. Aurrecoechea, C. et al, "A Review of Quality of Service Architectures," *ACM Multimedia Systems Journal*, November 1995.
6. Whitner R.B., “Designing Scaleable Applications Using CORBA,” *Integrated Network Management*, V, May 1997.
7. Maston M.C., “Using the World Wide Web and Java for Network Service Management,” *Integrated Network Management*, V, May 1997.
8. Orfali R. and D. Harkey, *Client/Server Programming with Java and CORBA*, Wiley Publication Co., 1997.
9. Campione, M. and K. Walrath, *The Java Tutorial: Object-Oriented Programming for the Internet*, Addison-Wesley, 1996.
10. Visibroker for Java by Visigenic Software, information available at *http://www.visigenic.com*
11. Coleman, D. et al, *Object-Oriented Development: The Fusion Method*, Prentice Hall Object Oriented Series, 1994.