NETWORK APPLICATIONS AND THEIR REQUIREMENTS

What are library networks used for right now and what will they be used for in two years? What level of performance do those applications require? Answer these questions before creating a network checklist. Analyze why that network is being built with a checklist that lists the components of a library network and specifies what type of technology to use. Only in this way can you justify why some technologies are better than others.

Base your analysis of the necessary applications on a consistent method of measuring performance requirements. The use of a common conceptual framework for describing performance requirements makes possible an understanding of the common or aggregate requirements applications place on the network. More importantly, you can identify the applications with the highest performance requirements. Their requirements serve as the critical measure of which technologies should be preferred in the checklist.

If a network can be built that meets these maximum requirements, other applications with lesser requirements will also perform well.

This section of the report describes library applications in terms of a common framework for measuring network performance requirements. The section begins by describing the interrelated set of concepts used to describe the performance requirements of network applications—quality of service, bandwidth, latency, and jitter—that provide the common framework for measuring application requirements. It then describes the two general classes into which the many library network applications may be divided and describes the requirements of each of these classes of applications in terms of this common framework.

The result of studying these applications' requirements is not a precise numerical estimate of needed bandwidth or other network requirements. An exact spreadsheet is not possible for a library with a specific number of users using certain applications to be able to calculate its exact bandwidth needs. The result is, for the most part, a qualitative rather than quantitative set of requirements. This analysis of applications provides a general characterization of how different types of applications depend on network performance, and how some applications require specific levels of bandwidth but others can operate with varying levels of bandwidth.

The analysis of library applications points out audio and video applications as the applications that place the heaviest requirements on library networks. Libraries should expect to implement audio and visual applications as part of their services and should use those applications' requirements as the measure of how a library network needs to be built. This will be the procedure used in this report's checklists.

Describing Network Application Requirements

Collectively, the performance of a network is referred to as its Quality of Service—abbreviated QoS in the networking literature. QoS

QoS—Quality of Service

Latency is the length of time for information to cross a network or a piece of that network.

Jitter is a variation in latency.

has two key measures: latency and jitter. Latency is another word for delay. It refers to the length of time for information to cross a network or a piece of that network. Jitter is variation in latency. A network with a large amount of jitter delivers information inconsistently, speeding up and slowing down unpredictably.

Latency is the product of multiple factors—the most obvious is bandwidth, which is the amount of information that can pass a point in a network in a given amount of time.

Bandwidth is traditionally measured in bits per second, for example, 100 Kbps. This is different from most computer applications where bytes are used as a measure of capacity. A bit is one-eighth of a byte. The use of bits rather than bytes to specify bandwidth is indicated by the use of a lower case b. Kbps indicates kilobits per second. Unfortunately, occasionally an upper case B will be used by mistake. If, in the context of a discussion of bandwidth or networking, an author uses a large B, for example, KBPS or KBps, it is safe to assume that bits not bytes are meant.

Insufficient bandwidth increases latency, because it causes information to take more time to pass through a network. Since bandwidth can have so much impact on latency, the bandwidth requirement of an application is a useful way to state the performance it requires from a network.

Bandwidth is not the only factor in the network that can cause increased latency. Network traffic must pass through networking devices on its way through the network. Each of these, from the computer used by the library patron to the devices used to connect networks to the remote server where the desired information is located, process the information transmitted over the network. The receipt of information, its processing, and then retransmission by each of these devices all add to the total network latency.

The delay in getting information through the network is also increased by the time needed to physically transmit information over any segment of the network. More bandwidth means more information can pass through a point in the network at the same time. Even if bandwidth was sufficient to simultaneously send all of a file through a network, physics would still limit how fast that information could travel across the network. An electrical signal only travels so fast down a cable no matter how great its bandwidth. Latency can result just from the distances spanned by a network.

Jitter, the variation in latency, occurs because the load on the network increases and decreases. When network load increases, network resources may be in contention. If a network goes from supporting a few users needing limited bandwidth to having a larger number of users demanding a lot of bandwidth, latency grows. Latency changes both because of the changing level of contention for the bandwidth available from the network and because of the changing amounts of information that must pass through network devices. Fluctuations between heavy and lighter network loads increase jitter.

Library Applications and their Network Requirements

The applications a network must support can be divided into two main categories based on whether they transmit information continually or in sporadic bursts. Most applications are bursty; they require sending a one-time transmission of information over the network. This is characteristic of Web browsing and most library applications. With bursty applications, users request information, that information transmits across the network, and then it is read or otherwise used. During the time it is being used, no additional transmission is occurring across the network.

In contrast, newer applications, such as audio and video, create a continuous flow of information across the network. While a person is viewing a movie or listening to a radio station over the Internet, the data from that movie or radio station streams, in that it continues to flow. The transmission of data is not a one-time event; it is ongoing.

This distinction is important because the network requirements of the two applications are different. Audio and video applications continually consume network resources. They are also more dependent on low-latency, high-bandwidth networks than are bursty applications, and are also more sensitive to jitter. Because of their greater demands on the network, audio and video applications set the highest level of requirements on library networks. The demands of these applications mostly guide the selection of technologies for inclusion in this report's networking checklists.

Bursty applications are described first below, followed by audio and video applications.

Audio and video applications set the maximum requirements on the network. Contrasting the two types of applications help make the greater demands of audio and video applications clearer. Discussing both also helps explain the collective impact of the bursty and audiovisual applications together on the same network.

Bursty is a one-time transfer of information.

Stream is most common with audio-visual applications or continuous flow of information across the network.

Bursty Applications

Bursty communication is a characteristic of both LAN (local area network) and Web applications. LAN applications include access to files on file servers or shared databases. They are typical office productivity applications that result from the need to share information within an organization. In libraries, LANs are also used to create access to shared information resources for the public. CD-ROMs with citation database and other resources may be mounted on a LAN. The LAN is also used as the source of software that may run on public or staff workstations.

Web applications take many forms. The Web can manage internal library communication. An example of an internal Web application is access to the library catalog's Web interface by a patron from a workstation within the library. Web applications most frequently provide access to resources outside the library. Examples include searching of library catalogs and citation databases, viewing the contents of digital image libraries, and accessing electronic full-text repositories.

Network Requirements of Bursty Applications

Bursty applications do not require ongoing synchronized delivery of content, so, unlike audio and video, no specific amount of bandwidth is required. The time it takes to download information largely depends on the bandwidth available to the network user. The more available bandwidth, the faster the information download. If a limited amount of bandwidth is available, quality of service degrades because the time to download information grows, but nevertheless downloading of material continues even at a slow pace.

To appreciate the impact of bandwidth on bursty applications, analyzing how long it would take a file of a specific size to be downloaded by a patron given different amounts of bandwidth is a useful exercise. This analysis explains the relation of bandwidth to the amount of time patrons spend waiting for information. The following table shows how long it would take a patron to download a 5 megabyte (MB) file given varying amounts of bandwidth:

Transmit times for a 5 Mb file.

Bandwidth	Time
56 Kbps	714 seconds
1 Mbps	40 seconds
10 Mbps	4 seconds
100 Mbps	0.4 second
1 Gbps	0.04 second

From 56Kbps to 1 Gbps, the amounts of bandwidth in this table represent the different amounts of bandwidth commonly available from different types of networks. A user experiences a clear difference in the delay and the quality of service when the amount of bandwidth is reduced. Experience with library patrons suggests whatever kind of information is involved, whether electronic journals or images, patrons want the network to give them at least the 4-second performance suggested above—and preferably the 0.4-second performance.

You can determine how much total bandwidth multiple users need by adding up the number of users and multiplying by a fixed amount of bandwidth per user. With bursty communications, you cannot identify an aggregate amount of bandwidth needed in a network by adding up the number of library computers used for Web browsing and the average size of the page accessed. Web users are not constantly receiving information; sometimes they are reading what they have obtained. With multiple users accessing the WWW, not all requests go out at the same time nor do all results come in at the same time. The number of users who can use a network connection before performance degrades is higher than might be suggested by dividing the amount of bandwidth available by a number of users and determining how much bandwidth should be available to each user.

Even though these kinds of calculations cannot produce an accurate estimate of the total bandwidth requirements of a network, they still have value. McClure (1995) calculated that 25 patrons using computers to access the Web, each consuming 56 Kbps, would require the bandwidth available from a 1 Mbps Internet connection. His estimate of the results of 25 patrons simultaneously using the network underlines the limited performance experienced by patrons during periods of heavy Internet use and when patrons simultaneously request large amounts of information. This can be typical of what happens in a library class when the teacher instructs the class to submit a search; all the students likely would do so at the same time. The industry says 100 users can share a 1 Mbps Internet connection. For bursty applications, the quality of service received for a given amount of bandwidth can vary dramatically.

In addition to requiring sufficient bandwidth to allow information to be downloaded with a low level of latency, bursty applications are also characterized by an asymmetric use of bandwidth. Different amounts of information are sent to and from other computers. For example, with Web applications, requests for Web pages are smaller than those pages. The request is usually little more than the URL for the page—an amount of text smaller than this paragraph. The page received is much larger. If the URL is the address for an electronic book or a journal article, the information sent back by the Web server might be many megabytes in size.

To summarize the quality of service requirements of bursty applications: they tend to be asymmetric, with more information received than sent. They are not susceptible to problems in jitter. Their major quality of service problem is latency. Sufficient latency is needed to make possible a quick enough download of information to satisfy library patrons. The amount of bandwidth available affects this speed. A low-latency, high-bandwidth network is needed to ensure library patrons can receive large files in a timely way. With bursty applications, many users can be supported simultaneously because they are not all using the network at the same time. At any given moment, only a few users may be downloading information; most users are using the downloaded information.

Streaming and Interactive Audio and Video Applications

There are two types of network audio and video applications. The first kind is streaming applications, such as RealVideo, that are not interactive. Information transmits just one-way; the user consumes information. The second are interactive audio and video applications, which provide two-way communication—the library patron both consumes and produces information. The streaming and interactive audio and video applications that use the Internet are presented below.

Audio

Of the multimedia Internet applications that have appeared in the last few years, one that has become the most widespread has been access to music through the Web. This has happened with the proliferation of Internet radio stations and downloadable MP3 files. Although the legality of much of this music distribution has been questioned, it has made online audio access widespread.

Although Internet radio involves streaming audio, many times MP3 files are not streamed over the Web; instead, they are first downloaded and then played by the recipient's computer.

Streaming audio applications developed on the Internet use compression mechanisms to limit the amount of bandwidth required for audio access. You can have fairly good receipt of audio over the Internet with a high-speed 56 Kbps modem connection. Higher-level, CD-quality audio requires more bandwidth—around 100 Kbps. This is not a huge amount of bandwidth, but because audio applications continually consume bandwidth, several audio users can quickly fill a lot of bandwidth.

Interactive audio applications have latency and jitter requirements that are more stringent than streaming audio. The amount of bandwidth required for a voice conversation is limited—64 Kbps is required by the standard phone system—but there are not the same opportunities for buffering interactive audio, as there are for streaming audio. Streaming applications such as Internet radio can store a large amount of audio data before it is delivered, thereby hiding the latency and jitter in the network from the user. With interactive applications, this is not possible as conversa-

Buffer is a data area shared by hardware devices or program processes that operate at different speeds or with different sets of priorities. (Definition compiled from www.whatis.com.)

tions may rapidly switch back and forth between the participants. Because the opportunity to buffer is less, interactive audio applications need lower latency.

Research (Siegel 2000) has found that voice latency can only grow so large before users become dissatisfied. The traditional telephone system, the public switched telephone network (PSTN), is an extremely high-performance system. Its latency is typically around 12 ms. Latency must remain less than about 300 ms. Jitter must also be limited, otherwise the speed of conversation slows down and speeds up.

VoIP (Voice over Internet Protocol) is the set of facilities that manage the delivery of voice information using Internet Protocol. (Definition compiled from www.whatis.com.)

TREND WATCH: *The most important interactive audio application will probably be Voice over IP (VoIP). A technology has been created to enable the Internet to be used for telephone calls. VoIP is being adopted by some organizations for their phone systems, such as Lakehead University. More and more parent organizations of special and academic libraries will probably adopt it as a way of placing phone calls. VoIP uses the same standard used for videoconferencing over the Internet, H.323. The software and hardware that implements H.323 can add significant latency to interactive audio communication. This delay can be over 100 ms (Siegel 2000). Combined with the latency created by the time needed to travel across the network using limited available bandwidth, you may find maintaining an adequate level of performance difficult. If latency from hardware and software processing is 100 ms and latency must be kept below 300 ms, little room exists for additional latency to be added by the network.*

Video

To learn more about streaming, see real.com.

Streaming video has many forms. Like streaming audio, these applications are sensitive to latency and jitter. The notable additional impact of supporting streaming video is the additional bandwidth needed to maintain an adequate level of latency. All forms of video require large amounts of bandwidth.

At the low end are technologies such as RealVideo where the video stream might only require hundreds of kilobits of bandwidth. Although this type of video supports the use of low-quality video streams, for example video with a 56 Kbps modem, RealVideo users have continually rising expectations. They have continued to demand higher-quality video from year to year. The author has been involved with Webcasting events with RealVideo for the last few years. Each year, the viewers have complained about video at last year’s level of quality and wanted higher levels of quality. The last event was broadcast at a maximum bandwidth of 512 Kbps.

This quality of video still does not get close to the quality of video library patrons expect from VHS tapes and television. Both media employ higher quality and larger pictures than RealVideo. Using MPEG-1 video encoding, video with the same level of quality as VHS tape requires 1.5 Mbps. Standard TV-quality video would require from 6 to 8 Mbps. The future standard of video quality to which library patrons will become accustomed, HDTV (high definition television) will require 30 to 40 Mbps. The amount of bandwidth required for video is striking. High-quality video requires many times more bandwidth than has been required for traditional library applications.

Interactive video has a different set of network requirements than streaming video. It requires bidirectional communications, but because videoconferencing users have been satisfied with lower-quality video for this purpose, its bandwidth requirements are not as great as MPEG-1 or MPEG-2 video. Videoconferencing can be done with 384 Kbps or less bandwidth. That 384 Kbps is regarded to be the minimal level of bandwidth that should be available to videoconferencing to have both acceptable voice and video performance, but it can still function at

128Kbps so long as the subjects do not move much. Videoconferencing functions at lower rates of bandwidth by letting the video quality degrade while preserving audio quality.

As videoconferencing can be thought of as an enhanced form of interactive audio, it has the same latency requirements as interactive audio. People have the same expectations for the performance of audio and video applications. Even if video quality degrades because of less than the desirable amount of bandwidth, as happens with 128 Kbps videoconferencing, people still expect the audio to perform within the same limited amount of latency as an interactive audio application.

Historically, videoconferencing has been accommodated by the implementation of special network lines for videoconferencing. The general bandwidth on a LAN or an Internet connection has not been used for this purpose. This is changing. The newest standard for videoconferencing, H.323, is designed to use the Internet, not special purpose networks, so it adds to the overall bandwidth requirements on the network.

Audio and Video Network Requirements

Library patrons want network audio and video applications to resemble those they have at home, with the low latency of the phone system, and the high video quality of television. This requires low latency networks and networks with sufficient bandwidth to ensure this low level of latency. Transmital of audio and video data requires more bandwidth. This bandwidth ranges from 64 Kbps for interactive audio up to 6 to 8 Mbps and higher for television-quality video. Unlike bursty applications, audio and video require this bandwidth on an ongoing basis. Each audiovisual application consumes a fixed amount of bandwidth that is not available to other users. The library network must be designed so users do not contend for network resources increasing latency and jitter.

Design Networks To Support Audio and Video

Some readers may object that their libraries will not be providing Internet-based audio and video services anytime soon, so this is not a relevant way to set network requirements. But the approach is valid because of the growth in the use of audio and video applications and because libraries cannot change some parts of their networks overnight. A sufficiently farsighted plan for development of a library network must account for current and future trends.

Some libraries have also recognized the opportunity to use audio and video. A number have created sites where videoconferencing can be done. Some have also started to include audio and video in their digital library projects, creating digital archives of these media.

The Library of Congress has digitized some films and put them on the Web. Several years ago, Indiana University created one of the first online digital music libraries. Many other libraries are exploring the potential for these technologies.

New multimedia applications are also becoming a more common experience for library patrons. Their personal familiarity with audio and video technologies is growing. Patrons bring their familiarity with audio

Links to Library of Congress: www.loc.gov

For an example of digital film delivery: www.memory.loc.gov/ammem/papr/sfhome.html

Link to Indiana Digital Music Library: www.dml.indiana.edu and video to the library and expect that at the library, like at home or work, audio and video can be used alongside other traditional media. The price of digital camcorders has continued to fall and computers are now often sold along with inexpensive PC video editing tools.

Library patrons are also increasingly exposed to audio and video on the Internet. The use of the Internet to distribute audio files has become widespread. This has not just happened through Napster and other exchanges of audio recordings. Many songs are also legitimately available on the Internet, having been put there by music publishers. Internet radio use has also become common. Likewise, more people are becoming familiar with Internet video materials. The Internet provides access to movie trailers, news videos, and other materials.

The growing availability of high-speed access to the Internet through DSL or Cable Internet services encourages the library patron's familiarity with multimedia. Increased availability of high-speed network connections both encourage Internet users to access audio and video materials over the Internet and encourage Web site developers to increase use of those media.

If changing parts of the library network every day were possible, planning for future needs wouldn't be necessary. Changing the network isn't easy; it can be expensive and time-consuming. Some parts of the network, for example the network wiring, can be expensive and might be replaced sooner than five or more years after it was installed.

Five years is an eternity in the world of technology. To plan this far ahead, expect to implement applications requiring more performance than any applications currently being considered. If your library won't need to support audio and video anytime soon, that is a strong recommendation to build with the less easily changed components of the network to meet those requirements.

A farsighted view of network requirements is also an asset in an environment where money for technical upgrades is not always available. You can make hardware choices able to cope with an unpredictable future that might bring new and more information-intensive applications into widespread use at a faster rate than expected. Don't plan too far out on the cutting edge in a budget-strapped institution. Those institutions tend to be risk averse and you have less of an opportunity to be wrong about a technical decision. Even within the scope of technologies that are not too risky, some technologies are more capable than others of meeting future needs. A focus on the future helps guide choices to these technologies.

Whether you believe your library will use audio and video as part of its library services in the near future, you should use these new media and their requirements as the means of evaluating what a network needs. Using such evaluations can help determine what investments are required to support audio and video, and when to make those investments at the most favorable moment. Preparing the network for the future puts the library in a position to meet the rising level of expectations library patrons have for access to information.

Conclusion

The advent of the multimedia age is placing new and greater demands on library networks, both on the library LAN and on the library's connection to the Internet. Libraries will be providers of access to audio and video information. Audio and video applications consume dramatically higher amounts of bandwidth than Web browsing or other bursty applications. They also consume bandwidth on a continual basis rather than allowing bandwidth to be shared by multiple users. Audio and video dedicate a large portion of bandwidth to one user on an ongoing basis, decreasing the number of library patrons who can be supported by a given amount of network bandwidth. Library networks require more bandwidth if they are to support these applications and simultaneously maintain a low latency for bursty applications. Networks also must be provisioned and designed to ensure limited latency and jitter for audio and video.

With an understanding of what audio and video applications require from the network, next you must use the following checklists that list what library LANs and the library's Internet services need to cope with these requirements.