Design and Implementation Guide-Part 1

December 1995

Jeff Baher
jbaher@cisco.com
Technical Marketing

Networked Multimedia

Overview

Networked multimedia applications are rapidly being deployed in campus LAN and WAN environments. From the corporate perspective, networked multimedia applications such as network TV or videoconferencing hold tremendous promise as the next generation of "productivity tools." But that's not all; the use of digital audio and video across corporate network infrastructures has tremendous potential for a slew of internal and external applications. The World Wide Web is the best example today of networked multimedia and its manifold capabilities.

Deploying networked multimedia applications in a campus LAN/WAN environment requires careful thought. The first consideration is how to enable sufficient bandwidth across the enterprise network for multimedia traffic. Next, if the multimedia applications support multicast transmission, is how to most effectively transmit the multicast traffic. This is particularly important to address in high-performance networks comprising virtual LANs (VLANs) and switches. Last is how to address quality of service (QoS) across the campus for those applications that require consistent guaranteed delivery.

As the leading router and switch manufacturer, Cisco Systems is also leading the way in building networked multimedia-enabled campus LANs and WANs. Through the use of Cisco's broad product offerings and the robust Cisco Internetwork Operating System (Cisco IOS™) software, Cisco offers a variety of different design scenarios that are optimized for networked multimedia applications.

The following document contains eight principal sections:

1 Introduction

2 Multimedia Basics: Digital Video and Audio

3 Networked Multimedia Applications

4 Campus LAN and WAN Bandwidth Solutions

5 Campus LAN and WAN Multicast Solutions

6 Campus LAN and WAN Quality of Service Solutions

7 Design Considerations for Networked Multimedia Environments

8 Conclusion and References

The intention of this document is to provide a solid understanding of networked multimedia and to highlight the critical issues that must be addressed when designing campus LAN and WAN environments. To this end, Section 7 ("Design Considerations for Networked Multimedia Environments") is crucial reading. This section is intended to bring together all the issues discussed throughout the document. In particular this section addresses how to build "smart" networks that will deliver both high performance and scalability for networked multimedia applications.

Cisco Confidential
Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

I. The Multimedia Revolution

Over the past few years, the concept of desktop, computer-based multimedia has gained considerable attention and acceptance. Some would even characterize it as a multimedia revolution. Gone are the days of dumb terminals, or even 286s or 68000, for the matter. Today its a world of CD-ROM equipped high-end 486s, Pentiums, 68040s and PowerPCs. In fact, more than eighty percent of personal computers sold during 1995 will be multimedia capable. This hardware revolution has, in turn, spawned a software revolution that has brought a wide-range of audio- and video-based applications to the desktop. It's not uncommon to see computers running video editing or image processing applications (Adobe Premier, Photoshop) in addition to basic "productivity" applications (word processing, spreadsheet and database applications).

The proliferation of more and more multimedia-enabled desktop machines has also spawned a new class of multimedia applications that operate in networked environments. These networked multimedia applications leverage the existing network infrastructure to deliver video and/or audio applications to end-users. Most notable are video conferencing and video server applications. With these application types, video and audio streams are transferred over the network either between peers or between clients and servers.

Many companies, for example, are turning to network multimedia applications to address employee training. Posed with the problem of having to fly entire sales forces back to headquarters for training, companies opt to use video-based distance learning applications to address the problem. With the distance learning application, the training session held at headquarters is broadcasted out to all remote sales offices over the existing LAN/WAN infrastructure. In this model, the sales force need only tune into the training session on his or her computer. Most companies will have to invest extra dollars into the network infrastructure, but this cost is offset by the expenses incurred to fly each member of the sales force over to headquarters.

Network multimedia is also changing the way people are communicating. Shared whiteboard applications offer users the ability to share a drawing space to hash out ideas while video conferencing applications allow users to see and hear people they are communicating with. The combination of both share whiteboard and video conferencing yields an effective vehicle for office communication.

Even the Internet is delivering network multimedia applications. RealAudio, for example, today offers real-time audio playback over the Internet. The quality isn't the best but give it time. A good place to go on the Internet for multimedia is the World Wide Web (aka the Web), the Internet's reservoir of information, graphics, sounds, etc. Here again the network multimedia revolution is evident. A wide variety of companies are turning to the Web to provide graphical front-ends (home pages) for public access. Internet surfers equipped with a Web browsing tool can access any of these Web home pages and gain access to a particular companies service. From a business perspective, the Web and its wide reach offer exciting new ways of conducting business. Check out Security First Bank, a new banking service offered on the Web. Using a new security method for data encryption, this bank allows customers to perform much of their banking transactions over the Web.

Put simply, network multimedia is revolutionizing the way people do business. And in the future network multimedia looks to only take on more. Especially with the Regional Bell Companies and Cable companies aggressively mapping out strategies for video-to-the-home, the future is very bright for wide spread deployment of network multimedia.

To successfully deliver networked multimedia it is important to understand both multimedia and networking. From a networking perspective, Cisco has identified three critical components to consider when deploying network multimedia applications in Campus LAN and WAN environments:

- Bandwidth: How much bandwidth do the applications demand and how much bandwidth can the network infrastructure provide
- Quality of Service: What level of quality of service does the application require and how can this be satisfied through the network
- Multicasting: Does the application utilize bandwidth saving multicasting techniques and how can multicasting effectively be supported across the network

Cisco has taken an active position in addressing the above three areas through new and innovative hardware platforms for routing, switching and ATM. From the bandwidth perspective, the Catalyst and Kalpana families of switching products, are key components in providing adequate bandwidth to networked multimedia clients and hosts. Cisco also has a variety of platforms that deliver high-speed network connectivity, namely Fast Ethernet, FDDI and ATM. Cisco has also developed a set of features integrated into the Cisco IOS (Internetworking Operating System), which runs on Cisco hardware platforms, to deliver consistent quality of service and to carry multicast traffic over the network.

This document intends to address the underpinnings of effectively deploying network multimedia applications. Specifically, this document will address:

- Multimedia basics: Analog video, digital video, video compression, digital audio
- Networked multimedia applications: Video conferencing, video broadcasts/multicasts, video servers, imaging, ATM-based video applications and their associated bandwidth requirements

Cisco Confidential
Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

- Bandwidth solutions for LAN and WAN environments: Switching, high-speed technologies, ATM, bandwidth-on-demand
- Quality of service: Priority queuing, Custom queuing, Weighted Fair Queuing, RSVP, ATM QoS
- Multicasting: IGMP, PIM, DVMRP (MBONE) Interoperability, SMRP, multicast application examples

II. Multimedia Basics

Much of today's video starts out as an analog signal, and thus a working knowledge of analog standards and formats is essential in understanding digital video and the digitization process.

There are four principal areas to understand when talking about analog video: broadcast standards, video signal standards, connectors and capture/storage formats. Below is a diagram showing how these four concepts interrelate.

Figure 1. Analog Video Concepts Overview

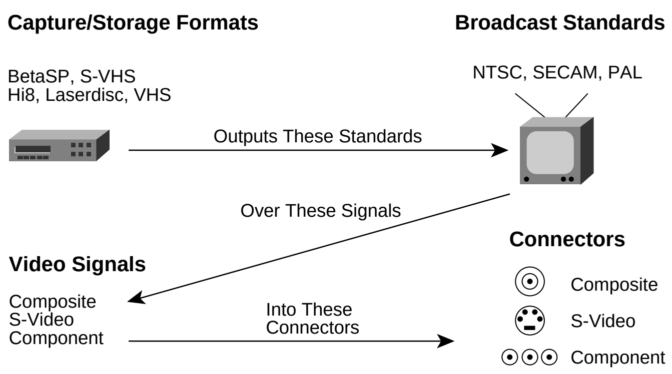

Broadcast Standards

There are three principal standards for analog broadcast transmission: National Television Standards Committee (NTSC), Phase Alternation Line (PAL) and Systemme Electronic Pour Couleur Avec Memoire (SECAM).

NTSC is the reigning standard for broadcast television in North and Central America and Japan. NTSC defines the screen with 525 vertical scan lines per frame and yields 30 frames per second. The scan lines refer to the number of lines from top to bottom on the television. The frames per second refer to the number of complete images or frames that are displayed per second. PAL dominates in Europe and is also found in the Middle East, Africa and South America. Lastly, SECAM is a PAL variant used in France, Russia and regions in Africa. Both PAL and SECAM standards are defined as 625 scan lines and refreshes the screen 25 times per second.

To produce an image on a television, an electron gun scans across the television screen from left to right moving from top to bottom.

Figure 2. Television Scan Gun Operation (Side and Front Views)

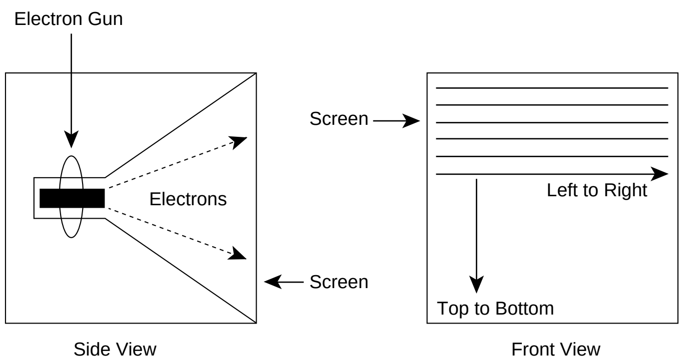

A problem arose with this technique because early television sets used a phosphor-coated tube. Consequently, by the time the gun finished scanning all the lines that the broadcast standard required, the lines at the top were starting to fade. To combat this, an interlaced technique was adopted whereby on the first pass from top to bottom, only every other line was scanned. With NTSC, this means that the first pass will generate 262 1/2 lines. The second pass will generate another 262 1/2 lines which are used to fill in the rest of the TV image.

Figure 3. Interlace Scan Process

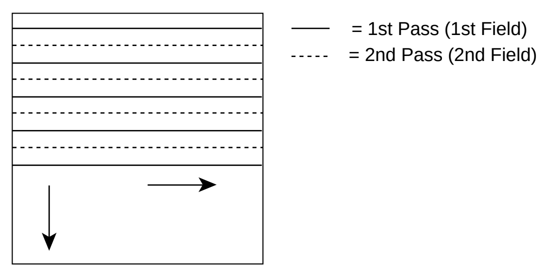

As the above diagram indicates, a frame represents the combination of the two passes or two fields. Hence in order for NTSC to deliver 30 frames per second, it must generate 60 fields per seconds.

As an aside, the 60 fields per second is derived from NTSC's clocking source. NTSC clocks it refresh intervals from the AC power. In the United States the AC power runs at 60 hertz or 60 oscillations per second. The 60 hertz yields the 60 fields per second with every two fields yielding a frame. In Europe the AC power clocks at 50 hertz. This yields 50 field per second or 25 frames per second.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Video Signals & Connectors

The video signals relate to how color and other video information is stored and transmitted. The transmission method also impacts the type of connector used on the video capture card when going from analog to digital. A composite signal, which carries all the information in one electrical channel, uses a one-hole jack called the RCA Phono connector. The S-Video signal, composed of two electrical channels, uses a four-pin connector called the mini-DIN connector. Finally, the Component connector signal uses three connectors. (See Figure 1 for diagrams.)

The Color Paradigm

In 1953, the NTSC had a problem: Black and white televisions were pervasive, but color television sets were coming. The NTSC had to devise a standard that would be backward compatible (i.e., the broadcast signal would yield black-and-white reception on the older sets and color on the new color sets).

Black-and-white televisions receive one signal called luminance, in video terms referred to as the "Y" signal. Each screen pixel is defined as some range of intensity between white (total intensity) and black (no intensity). To maintain compatibility with older black and white sets, the NTSC had to set a color standard that kept luminance signal separate and also provided color information required for newer color television sets.

In the digital world, colors are typically expressed using Red, Green and Blue (RGB). The analog world has also embraced the RGB standard, at least on the acquisition side, where most cameras break the analog signal into RGB components.

Unfortunately, the NTSC could not use RGB as the color television standard because the old black and white sets could not decode RGB signals. Instead, they had to send a luminance signal for black-and-white sets and fill in the color information with other signals, called hue and saturation, or U and V. For this reason, where the digital world uses RGB, the analog world, especially broadcast television uses YUV.

Component, S-Video and Composite

Figure 4 traces an analog signal from RGB to composite. On the far left is the RGB capture in which storage channels are maintained for each of the three primary colors. RGB, however, is an inefficient analog video storage format for two reasons. First, to use RGB, all three color signals must have equal bandwidth in the system, which is often inefficient from a system design perspective. Second, since each pixel is the sum of Red, Green and Blue values, modifying the pixel forces an adjustment of all three values. In contrast, when images are stored as luminance and color formats (i.e. YUV format), a pixel can be altered by modifying only one value.

Component video means that separate channels are maintained for each color value, both in the recording device and the storage medium. This minimizes noise that occurs when two signals are combined in one channel and thus delivers the highest fidelity.

After NTSC encoding, the hue and saturation channels (U&V) are combined into one chrominance channel, the C channel. A video signal called the S-Video carries separate channels for the luminance (Y) and the chrominance (C) signals. S-Video is also called Y/C video.

To play on old black-and-white televisions, ultimately all color and other information must be combined into one YUV channel, a composite signal. Technically, a composite signal is any signal that contains all the information necessary to play video. In contrast, any one individual channel of component or Y/C video would not be sufficient to play the video.

Figure 4. RGB to NTSC Encoding

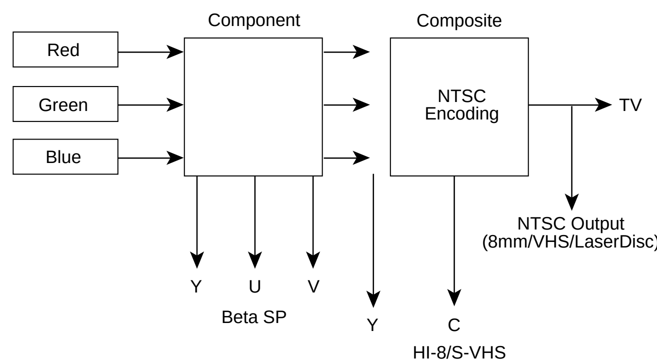

Capture and Storage Formats

The first major difference between capture formats is how the color information is stored. BetaSP stores the color information as component video with three separate channels. Hi-8 and SVHS uses the two-channel Y/C video, while VHS, 8mm and laserdiscs store all color information in one channel. As for quality, the composite signal provides the lowest quality because all signals are combined which in turn has the highest potential for noise. The S-Video signal produces less noise since the two signals are isolated in separate channels, not merged together. As expected, the component signals provides the highest quality signal, since all components are maintained in separate channels. Remember, from the capture board perspective, the image quality at best is only as good as the signal that it can accept. Refer to Table 1 (Analog Storage Formats) for a matrix of analog capture and storage standards.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Table 1. Analog Storage Formats

	BetaSP	SVHS/Hi-8	VHS/8mm	Laserdisc
Color signal	Component	S-Video	Composite	Composite
Lines of resolution	750	400	200	400
Signal to noise (s/n) ratio	50db	47db	47db	47db

Lines of Resolution

As Table 1 indicates, each storage format delivers different lines or resolution. Resolution is a measure of an image's quality. From the viewer's perspective, an image with higher resolution will yield a sharper picture quality than a lower resolution image.

Most consumer televisions display roughly 330 lines of horizontal resolution. Typically in broadcast environments, high-end cameras will be used to capture video. With these cameras and their associated storage formats, they can deliver horizontal resolutions in the neighborhood of 700 lines. By recording in high-resolution, this means that multiple generations of the video can be made without a noticeable difference from the users perspective. Each time a copy is made, the original image loses some of its resolution. Starting with a high resolution allows a lot of room for impairments before affecting the end result.

It follows that starting with a lower resolution image will mean there is less room to manipulate the image before the viewer notices its effects.

Digitizing Video

Digitizing video involves taking an analog video signal and converting it to a digital video stream using a video capture board (see figure 5). Today, PC, Macintosh, and UNIX workstations offer video capture capabilities. In some cases, though, the capture equipment will be a third-party add-on. The analog video source can reside on any of the above video storage formats (BetaSP, SVHS, laserdisc or VHS) or it can be a live video feed from a camera. The source can be connected to the capture card using any of the above three connectors: component, S-video or composite (depending of course on what connector support the capture card offers).

When capturing and digitizing video there are three critical components to keep in mind:

- Resolution
- Color depth
- Frame rate

The *resolution* refers to the horizontal and vertical dimensions of the video session. A full screen video session is typically 640 horizontal pixels x 480 verticals pixels. Full screen video uses these dimensions (640x480) because it yields the 4:3 aspect ratio of standard television. As discussed earlier, NTSC specifies 525 vertical scan lines of which 483 lines are used to display video. The balance is used for signaling and is referred to as the vertical blanking interval. If a television set rolls dues to a vertical hold problem, a black bar is visible. This black bar is comprised of the extra 43 scan lines. Note that since the NTSC image uses 483 vertical lines, capturing at 640 by 480 will mean that 3 lines will be dropped during the digitization process, a negligible difference.

The *color depth* refers to how many bits are used to express color. At the high-end is 24-bit color which is capable of displaying 16.7 million colors. The 24-bit color (or 3-byte color) is the aggregate of 8-bits of Red plus 8-bits of Green plus 8-bits Blue. The 8 bits are used to express color intensity from 0 to 255. Other common color depths are 16-bit and 8-bit, which yield roughly 65000 and 256 colors, respectively using color palettes.

Finally, the *frame rate* refers to how many frames are displayed per second. To deliver NTSC quality video, 30 frames per second will be displayed. PAL and SECAM environments would uses 25 frames per second. And for those interested in creating a digital cinema, motion pictures run at 24 frames per second.

Based on the above three criteria, it is a simple mathematical operation to determine how much bandwidth a particular video stream will require.

To carry uncompressed NTSC-quality digitized video, for example, approximately 27 Megabytes per second are necessary in order to deliver the video over the network. This number is derived from the following calculation:

> 640 x 480 (screen resolution) x 3 (# of bytes of color) = bandwidth per frame
>
> bandwidth per frame x 30 (number of frames per second) = bandwidth per second
>
> 640 x 480 x 3 x 30 = 27.648 Mbytes or 221.184 Mbits per second

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 5. Analog to Digital Video Conversion

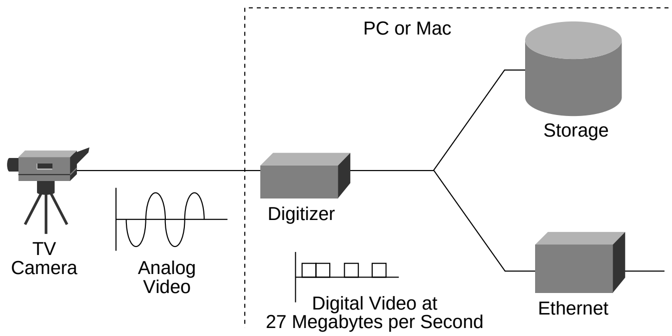

Full screen NTSC-video quality raw bandwidth requirements are shown in the figure below.

Figure 6. Full screen Video Raw Bandwidth Requirements

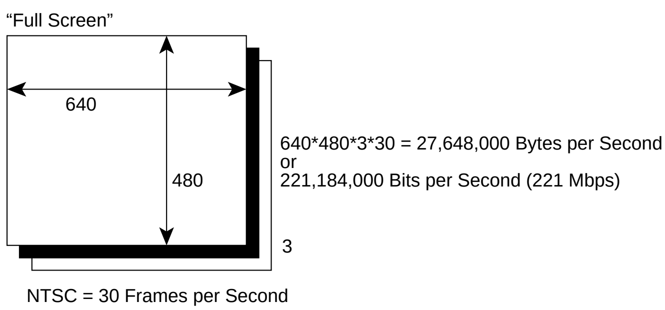

As the calculation above indicates, full-motion, full-color digital video requires considerably more bandwidth than today's typical packet-based network can support. Fortunately there are two techniques that can be employed to cut down on the bandwidth consumption:

- Modify video capture parameters
- Use video compression technology

Modifying Video Capture Parameters

Modifying video capture parameters involves manipulating any of the above three capture components: resolution, color depth or frame rate. To reduce bandwidth consumption, it is not uncommon to modify all three variables. For example, some multimedia applications capture video at 320 x 240 with 8-bit color and at a frame rate of 15 frames per second. With these parameters, bandwidth requirements drop to 1.152 Mbytes per second or 9.216 Mbits per second. While this level of bandwidth would stress an Ethernet network, it is no problem for 16MB Token-Ring, not to mention the 100MB technologies.

The other common approach to reducing bandwidth requirements of digital video is to use a video compression technology. Video compression is performed using a CODEC (Coder/Decoder or Compressor/Decompressor). The CODEC, which can be implemented either in software or hardware is responsible for taking a digital video stream and compressing it or receiving a pre-compressed video stream and decompressing it. Most PC, Macintosh and UNIX products combine the digitizer with the CODEC but remember these are separate processes.

Video Compression

Video compression is a process whereby a collection of algorithms and techniques replace the original pixel-related information with more compact mathematical descriptions. Decompression is the reverse process of decoding the mathematical descriptions back to pixels for ultimate display. At its best, video compression is transparent, or invisible to the end user. The true measure of a video CODEC is how little the end-user notices its presence, or how effectively it can reduce video data rates without adversely affecting video quality.

Figure 7. Post Digitization Video Compression

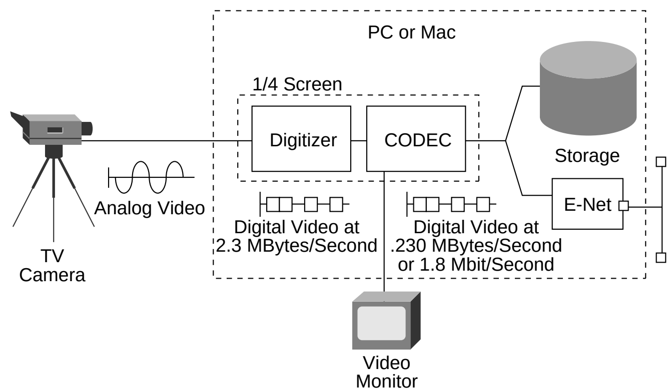

When discussing compression its is important to make a distinction between *lossy* and *lossless* compression techniques. Lossless compression techniques create compressed files that decompress into exactly the same file as the original, bit for bit. Lossless compression is typically used for executables (applications) and data files where any change in digital make-up renders the file useless. Lossless compression is used by products such as STAC and DoubleSpace to transparently expand hard drive capacity, and products like PKZIP to pack more data onto floppy drives. STAC and Predictor are supported in Cisco IOS for data compression over analog and digital circuits.

In general terms, lossless compression techniques identify and utilize patterns within files to describe the content more efficiently. This works quite well for files with significant redundancy, such as database or spreadsheet files. As a whole, however, these techniques typically yield only about 2:1 compression which barely dents high-resolution uncompressed video files.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Lossy compression, used primarily on still image and video image files, creates compressed files that decompress into images that look similar to the original but are different in digital make up. This “loss” enables such techniques to deliver from 2:1 to 300:1 compression. Run an executable through a lossy compressor and there will be serious problems. A 24-bit image through the same compressor, on the other hand, might have a few changed pixels or altered color shades, but overall the image will experience little change to the human eye. The effects of lossy compression is further minimized with video because each image is displayed for only a fraction of a second (1/15 or 1/30 or a second depending on the frame rate).

There are a wide range of lossy compression techniques available for digital video. A simple rule, however, is that the higher the compression ratio, the greater the loss. And as the loss increases so too do the number of artifacts. An *artifact* refers to a portion of a video image that is noticeably lossy (i.e. there is little or no information for the pixels in the affected area).

Interframe Compression

Video compression uses two basic compression techniques: interframe compression (also known as temporal compression) and intraframe compression (also known as spatial compression), which occurs within individual frames. Interframe compression, uses a system of key and delta frames to eliminate redundant information between frames. Key frames store a complete image, while delta frames, or difference frames, record only interframe changes. Interframe compression is also referred to as temporal compression because the compression is applied along the time dimension.

Depending on the video compression algorithm, interframe (temporal) and intraframe (spatial) may be employed. MPEG (Motion Picture Experts Group), for example, compresses both spatially (using JPEG [Joint Photographic Experts Group]) and temporally using a separate algorithm. Motion-JPEG, on the hand, only uses spatial compression.

As mentioned above, interframe compression uses a system of key and delta frames to eliminate redundant information between frames. In some schemes, the key frames will be compressed and in other cases they won’t. Either way, the key frames serve as a reference source for delta frames. Delta frames contain only pixels that are different from either the key frame or the immediately preceding delta frame, whichever frame they reference during compression. During decompression, delta frames look back to their respective reference frames to fill in missing information.

Different compression techniques use different sequences of key and delta frames. For example, most Video for Windows® CODECs, calculate interframe differences between sequential delta frames during compression. In this case, only the first delta frame relates back to the key frame. The rest of the frames relate back to the immediately preceding delta frame. Other compression schemes will have all delta frames relate back to the key frame, like MPEG.

Whatever the technique, all intraframe compression techniques derive their effectiveness from interframe redundancy. Low-motion video sequences such as the “talking head” (essentially just the head and shoulders of a person), have a high-degree of redundancy, which limits the amount of compression required to reduce the video to the target bandwidth.

Up until now, interframe compression has addressed only pixel blocks that remained static between the delta and reference frame. Some CODECs, however, increase compression by tracking moving blocks of pixels from frame to frame. This is called motion compensation, or dynamic carryforwards, since the data carried forward from key frames is dynamic not static. Consider a video clip where a person is waving their arm. If only static pixels were tracked between frames, no interframe compression would occur with respect to the moving parts of the person simply because those parts are not located in the same pixel blocks in both frames. If the CODEC, however, could track the motion of the arm, the delta frame description could tell the decompressor to look for particular moving parts in other pixel blocks, essentially tracking the moving part as it moves from one pixel block to another.

While dynamic carryforwards are helpful they are not always implemented. In many cases it comes down to the fact that the capture board cannot scale resolution and frame rate, digitize, and hunt for dynamic carryforwards at the same time.

Dynamic carryforwards typically mark the dividing line between hardware and software based CODECs. Hardware CODECs, as the name implies, are usually add-on boards that provide hardware compression and decompression operations. The benefit from hardware CODECs is that they do not place any additional burden on the host CPU in order to execute video compression and decompression.

Software CODECs, on the other hand, are executed using the host CPU and typically require no additional hardware. The benefit here is that the CODEC is typically less expensive and installation tends to be easier. In fact it can even be done over the network from a central point. Since the software CODECs use the host’s CPU to perform the compression and decompression often times they are limited by the host CPU’s processing capabilities. This in turn limits the CODEC’s functionality preventing certain techniques such as advanced tracking schemes.

Cisco Confidential
Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Intraframe Compression

Intraframe (spatial) compression is performed solely with reference to information within a particular frame. It's performed on pixels in delta frames that remain after interframe compression and on key frames.

While intraframe techniques are often the most hyped, overall CODEC performance relates more to interframe efficiency than intraframe efficiency. Below is a discussion of the principal intraframe compression techniques.

RLE

Run Length Encoding, RLE, is a simple lossless technique originally designed for data compression and later modified for facsimile. RLE essentially encodes "runs" of different pixel lengths. Essentially RLE takes an image an compresses it based on "runs" of pixels. Unfortunately, most 24-bit real-world videos, don't have long runs of identically colored pixels. While RLE works well in black-and-white facsimile world, it is less efficient for color video.

JPEG

JPEG has been adopted as a standard by two international standards organizations, the ITU (formerly CCITT) and the ISO. JPEG is most prevalent in still image compression. JPEG works first by encoding the video data. Here data is converted into frequency space for the discrete cosign transform (DCT) analysis. DCT starts by dividing the image into 8x8 blocks, then converts the colors and pixels into frequency space by describing each block in terms of the number of color shifts (frequency) and the extent of the change (amplitude).

Because most natural images are relatively smooth, the changes that occur most often, or high-frequency changes, have low amplitude values, meaning that the change is minor. In other words, images that have many subtle shifts among similar colors, but few dramatic shifts between very different colors.

Next, quantization and amplitude values are categorized by frequency and averaged. This is the lossy stage because the original values are permanently discarded. However, because most of the picture is categorized in the high-frequency/low-amplitude range, most of the loss occurs among subtle shifts that are largely indistinguishable to the human eye.

After quantization, the values are further compressed through RLE using a special zigzag pattern designed to optimize compression of like regions within the image. At extremely high compression ratios, more high-frequency/low-amplitude changes get averaged, which can cause an entire pixel block to adopt the same color. This causes a blockiness artifact that's characteristic of JPEG-compressed images. JPEG is used as the intraframe technique for MPEG.

Vector Quantization

Vector quantization (VQ) is similar to JPEG in that it divides the image into 8x8 blocks. The difference between VQ and JPEG has to do with the quantization process. VQ is a recursive, or multistep, algorithm with inherently self-correcting features.

With VQ, similar blocks are categorized and a reference block is constructed for each category. The original blocks are then discarded. During decompression, the single reference block will replace all of the original blocks in the category.

After selecting the first set of reference blocks, the image is decompressed. Comparing the decompressed image to the original will reveal many differences. To address the differences, and additional set of reference blocks are created that fill in the gaps created during the first estimation. This is the self-correcting aspect of the algorithm. The above process is repeated to find a third set of reference blocks to fill in the remaining gaps. These reference blocks are all posted in a lookup table to be used during decompression. The final step is to use lossless techniques such as RLE to further compress the remaining information.

VQ compression is by its nature computationally intensive. However, decompression, which simply involves pulling values from the lookup table is simple and fast. VQ is a public-domain algorithm used as the intraframe technique for both Cinepak and Indeo.

The discussion above provides the foundation for a variety of hardware and software based video compression standards. In many cases the network multimedia application will dictate the video compression algorithm used. For example, Intel's ProShare Video Conferencing uses the Indeo standard or H.261, while Insoft Communique! uses Cell B compression. With some applications, the end-user can specify the compression algorithm.

End-User Video Compression Algorithms

Below are brief descriptions of the more popular end-user video compression algorithms. Note that some algorithms require dedicated hardware while others can use a combination of software and host CPU cycles.

- *MPEG1:* MPEG1 defines a bit stream for compressed video and audio optimized to fit into a bandwidth of 1.5Mbits per second. This rate is special because it is the data rate of uncompressed audio CDs and DATs. Typically MPEG1 is compressed in non-real time and decompressed in real time. MPEG1 compression is typically performed with hardware while decompression can be done with either hardware or software.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

- *MPEG2:* MPEG2 is intended for higher quality video on demand applications for products like the "set top box." MPEG2 runs at data rates between 4 and 9 Mbps. MPEG-2 and variants are considered by Regional Bell carriers and Cables companies for delivering video-on-demand to the house as well as for delivering HDTV broadcasts. There are currently MPEG-2 chipsets on the market that perform real-time encoding. There are also real-time MPEG2 decompression boards available. Additionally, there is a specification for MPEG2 adaptation over ATM AAL5.
- *MPEG4:* MPEG4 is a low bit rate compression algorithm intended for 64Kbps connections. MPEG4 is targeted at a wide-range of applications including mobile audio and visual applications and electronic newspaper sources.

  For a listing of MPEG-based products, refer to http://www.crs4.it/rluigi/mpeg/mpegcompanies.html
- *M-JPEG:* M-JPEG, or Motion-JPEG, can be implemented in either hardware or software. M-JPEG is the aggregation of a series of JPEG-compressed images.
- *Cell B:* CellB is part a family of compression techniques developed by Sun Microsystems. CellB is designed for real-time applications such as video conferencing which requires real-time video transmission. CellB's counterpart, CellA, is intended for non-real time applications where encoding does not need to take place in real-time. Both CellA and CellB use VQ and RLE to achieve compression.
- *Indeo:* An offshoot of DVI, Digital Video Interactive, Indeo was developed by Intel. Indeo uses VQ as its intraframe engine. Intel has release three versions of Indeo. Version 2.1 was a strictly intraframe technology, largely focused on Intel's popular capture board, the Smart Video Recorder. In late 1993, Intel introduced Indeo 3.1, which incorporated interframe compression as well. Indeo 3.2 is now available. Indeo requires a hardware add-on for video compression but decompression can take place in software on a high-end 486 or Pentium.
- *Cinepak:* Cinepak was developed by SuperMatch, a division of SuperMac Technologies. Cinepak was first introduced as a Macintosh CODEC and migrated to the Windows platform in 1993. Cinepak debuted in Video for Windows release 1.1. Like Indeo, Cinepak uses VQ as its intraframe engine. Cinepak offers the widest cross-platform support of all the CODECs. Cinepak is also available for 3D0, Nintendo and Atari platforms.
- *Apple Video:* Apple Video is the compression technique used with applications like Apple's QuickTime Conferencing.
- *H.261:* H.261 is the compression standard specified under the H.320 Video Conferencing standard[1]. H.261 describes the video coding and decoding methods for the moving picture component of audiovisual services at the rate of p x 64 kbit/s, where p is in the range 1 to 30. It describes the video source coder, the video multiplex coder and the transmission coder.

  H.261 defines two picture formats: CIF (Common Intermediate Format) has 288 lines by 360 pixels/line of luminance information and 144 x 180 of chrominance information; and QCIF (Quarter Common Intermediate Format) which is 144 lines by 180 pixels/line of luminance and 72 x 90 of chrominance. The choice of CIF or QCIF depends on available channel capacity—e.g. QCIF is normally used if p<3.

  The actual encoding algorithm is similar to (but incompatible with) that of MPEG. Another difference is that H.261 needs substantially less CPU power for real-time encoding than MPEG. The algorithm includes a mechanism which optimizes bandwidth usage by trading picture quality against motion, so that a quickly-changing picture will have a lower quality than a relatively static picture. H.261 used in this way is thus a constant-bit-rate encoding rather than a constant-quality, variable-bit-rate encoding.

Hardware versus Software CODECs

As a rule of thumb, the more CPU cycles given to video compression and decompression, the better the performance. This can be achieved either by running software CODECs on fast CPUs (Pentiums, PowerPCs, RISC) or by investing more money in dedicated hardware add-ons such as an MPEG playback board. Keep in mind that in some cases the application that is being used will dictate hardware or software compression and decompression. Insoft's INTV video multicast package, for instance, uses a hardware-based compressor in the UNIX workstation, but uses a software-based decompressor for the PC clients. The implication is that in order to run with this product, the PCs may need to be upgraded in order to deliver the requisite processing capabilities.

Compression Ratios

Any of the above compression standards are helpful in reducing the amount of bandwidth needed to transmit digital video. In fact, today digital video can be compressed up to 20:1 and still deliver a VHS-quality picture.

Below is a table that shows digital video compression ratios and the approximate quality that they yield vis a vis analog video formats.

1. Cisco recognizes the importance of H.320 and its related protocols (H.323 and H.324) for videoconferencing interoperability. Currently, however, Cisco has not committed a timeframe for supporting these protocols.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Table 2. Image Quality as a Function of Compression Ratio

Video Compression Ratio	Analog Picture Quality Equivalent
20:1	VHS
10:1	S-VHS/Hi-8
4:1	Broadcast quality

As the table above indicates, fairly high ratios of video compression can be used while still preserving a high-quality video image. This in turn is great from a networking perspective simply because the amount of bandwidth required can be dramatically reduced without sacrificing too much of the end result. A typical MPEG1 video stream (640x480, 30 FPS), for example, runs at roughly 1.5Mbits per second.

Digital Audio

Like digital video, digital audio in many cases starts off as an analog source. As a result an A to D (analog to digital) conversion must be made. Converting an analog signal to a digital signal involves taking a series of samples of the analog source. The aggregation of the samples yields the digital equivalent of the analog sound wave.

Below are two diagrams each showing an analog sound wave. Figure 8 shows a low sampling rate while Figure 9 shows a high sampling rate.

Figure 8. Low Sampling Rate

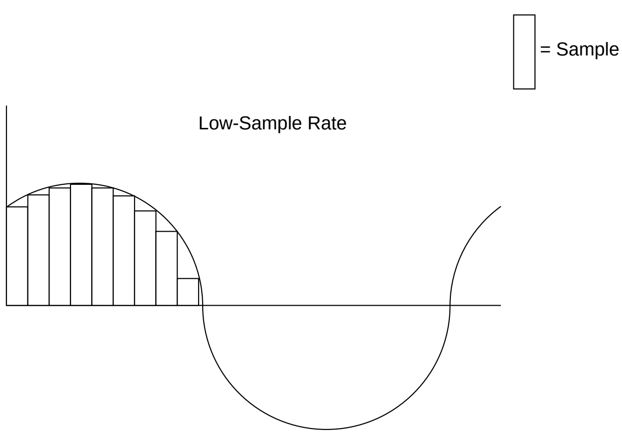

Figure 9. Low Sampling Rate

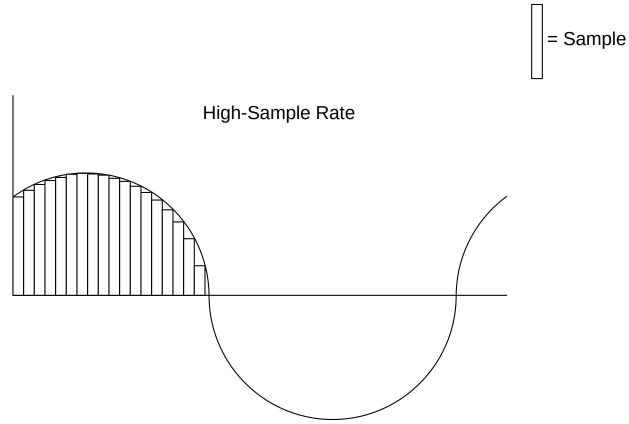

While in both cases, the analog signal has been converted to a digital signal, the higher sampling rate will deliver higher quality because it has more reference points to replicate the analog signal.

The sampling rate, or the number of samples per second of the analog source, is one of three criteria that determine the quality of the digital version. The other two determining factors are the number of bits used to represent each sample and the number of channels.

Sampling rates are often quoted Hertz (Hz) or Kilohertz (kHz). Sampling rates are always measured per channel, so for stereo data recorded at 8,000 samples per second (8 kHz), there would actually be 16,000 samples per second.

The table below shows common sampling rates.

Table 3. Common Sampling Rates

Sample/sec	Description
8000	Exactly 8000 samples/sec is a telephony standard that goes together with u-law encoding.
11 k	Either 11025, a quarter of the CD sampling rate, or half the Mac sampling rate (perhaps the most popular rate on the Mac).
16000	Used by the G.722 compression standard.
18.9 k	CD-ROM/XA standard.
22 k	Either 22050, half the CD sampling rate, or the Mac rate; the latter is precisely 22254.545454545454 but usually misquoted as 22000. (Historical note: 22254.5454... was the horizontal scan rate of the original 128k Mac.)

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Table 3. Common Sampling Rates (Continued)

Sample/sec	Description
32000	Used in digital radio, NICAM (Nearly Instantaneous Compandable Audio Matrix [IBA/BREMA/BBC]) and other TV work, at least in the UK; also long play DAT and Japanese HDTV.
37.8 k	CD-ROM/XA standard for higher quality.
44056	This weird rate is used by professional audio equipment to fit an integral number of samples in a video frame.
44100	The CD sampling rate. DAT players recording digitally from CD also use this rate.
48000	The DAT (Digital Audio Tape) sampling rate

Many of today's multimedia applications include audio support. Some applications include hardware for digitizing audio while other applications will rely on third-party add-ons for audio support. Check with the application vendor to learn how audio is handled.

There is an emerging tendency to standardize on only a few sampling rates and encoding styles, even if the file formats may differ. The suggested rates and styles are:

Rate (samp/sec) style mono/stereo
8000 8-bit U-LAW mono
22050 8-bit linear unsigned mono and stereo
44100 16-bit linear signed mono and stereo

Audio Compression

Strange though it seems, audio data is remarkably hard to compress effectively. For 8-bit data, a Huffman encoding of the deltas between successive samples is relatively successful. For 16-bit data, companies like Sony and Philips have spent millions to develop proprietary schemes. Information about PASC (Philips' scheme) can be found in *Advanced Digital Audio* by Ken C. Pohlmann.

Public standards for voice compression are slowly gaining popularity, e.g. CCITT G.721 (ADPCM at 32 kbits/sec) and G.723 (ADPCM at 24 and 40 kbits/sec). (ADPCM == Adaptive Delta Pulse Code Modulation.)

GSM 06.10 is a speech encoding in use in Europe that compresses 160 13-bit samples into 260 bits (or 33 bytes), i.e. 1650 bytes/sec (at 8000 samples/sec). There are also two US federal standards, 1016 (Code excited linear prediction (CELP), 4800 bits/s) and 1015 (LPC-10E, 2400 bits/s).

Apple has an Audio Compression/Expansion scheme called ACE (on the GS) MACE (on the Macintosh). It's a lossy scheme that attempts to predict where the wave will go on the next sample. There's very little quality change on 8:4 compression, somewhat more for 8:3. It does guarantee exactly 50% or 62.5% compression, though.

Digital Audio Hardware

The table below shows some of the more popular sound cards for PC, Mac, and UNIX.

Table 4. Popular Sound Cards for PC, Mac, and UNIX

Machine	Bits	Max Sampling Rate	Number of Output Channels
Mac (all types)	8	22k	1
Mac (newer ones)	16	64k	4
Apple IIgs	8	32k / >70k	16(st)
PC/soundblaster pro	8	22k st, 44.1k mo	1(st)
PC/soundblaster 16	16	44.1k	1(st)
PC/turtle beach multisound	16	44.1k	1(st)
PC/cards with aria chipset	16	44.1k	1(st)
PC/roland rap-10	16	44.1k	1(st)
PC/gravis ultrasound	8/16	44.1k	14-32(st)
Sun SPARC	U-LAW	8k	1
Sun SPARC. 10	U-LAW, 8, 16	48k	1(st)
NeXT	U-LAW, 8, 16	44.1k	1(st)
SGI Indigo	8,16	48k	4(st)
SGI Indigo2, Indy	8,16	48k	16(st, 4-channel)
HP9000/705,710, 425e	U, A-LAW, 16	8k	1
HP9000/715,725, 735	U, A-LAW, 16	48k	1(st)

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Table 4. Popular Sound Cards for PC, Mac, and UNIX (Continued)

Machine	Bits	Max Sampling Rate	Number of Output Channels
HP9000/755 option:	U, A-LAW, 16	48k	1(st)
NCD MCX terminal	U, A, 8, 16	52k	1(st)

III. Network Multimedia Applications

There are a wide-range of networked multimedia applications to choose from so it is important to understand upfront what the intended goals are of deploying a networked multimedia application. Additionally, it is important to understand the bandwidth implications of the chosen application.

Today's network multimedia applications can be broken down into the following four categories:

- Point-to-Point, bi-directional
- Point to Multipoint, bi-directional
- Point-to-Point, unidirectional
- Point to Multipoint, unidirectional

Point-to-Point, bi-directional, as its name implies, refers to a class of applications that deliver real-time point-to-point communication. The key here is that the process is bi-directional meaning that video is transmitted in both directions in real time.

Point-to-Multipoint, bi-directional, is similar to the point-to-point model, except this model employs multiple video senders and receivers. In this model multiple clients can be sending and receiving video stream in real time.

Point-to-Point, unidirectional, refers to a point-to-point communication in which video is transmitted in only one direction. The video itself can be either stored video streams or real-time streams from a video recording source.

Point-to-Multipoint, unidirectional, is similar to its point-to-point counterpart except in this case the video is transmitted to a group of clients not just one. The video is still unidirectional, however, no there are more recipients. Again the video can come from a storage device or a recording source.

Point-to-Point, Bi-directional Applications

Figure 10.

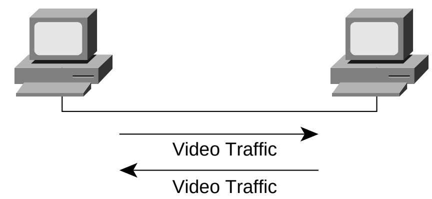

- Audio and/or Video conferencing
- Shared whiteboard
- Application sharing

Audio and Video conferencing applications provide a real-time interactive environment to two end-users. Often times these applications will also include a shared whiteboard application or with application sharing functionality. Share whiteboard applications provide a common "whiteboard" both end users can see and draw on. Shared whiteboards, also referred to as collaborative workspaces, are particularly useful in conversations where a picture is truly worth a thousand words. Application sharing is also a useful and productive tool. With application sharing, a local user can launch an application (Microsoft Access, for example) and the remote user can view and work with it as though the application were installed locally. Now co-workers at opposite ends of a network can collaborate using a shared application regardless of where the application actually resides.

Point-to-Multipoint, Bi-directional Applications

Figure 11.

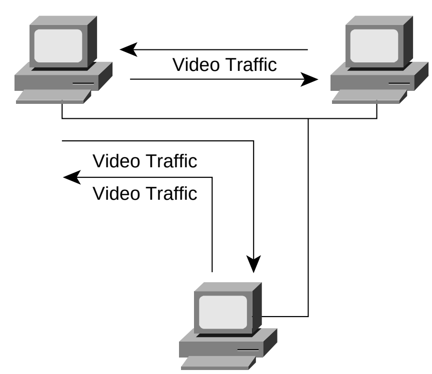

- Interactive video
- Video conferencing

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Interactive video, such as video kiosks, deliver video to multiple recipients. The recipients can interact with the video session controlling start and stop functions. The video content can also be manipulated by end-user interaction. Some kiosks, for example, have a touch pad which will deliver different video streams based on the pad selected.

As discussed earlier, video conferencing provides end-user video communication across the LAN or WAN. Like a telephone call where multiple listeners can be "conferenced" in, the same can be done with certain video conferencing applications. A three way video conference call, for example, can occur in which each person can receive video and audio from the other two participants.

Point-to-Point, Unidirectional

Figure 12.

- Video server applications
- Multimedia-enable e-mail applications

In the above two applications, video clips are stored centrally and fed to a remote user as needed where they are then decompressed for viewing. The information is stored centrally pre-compressed. The end-use initiates the viewing process by downloading the stream across the network to the video decompressor.

Point-to-Multipoint, Unidirectional

Figure 13.

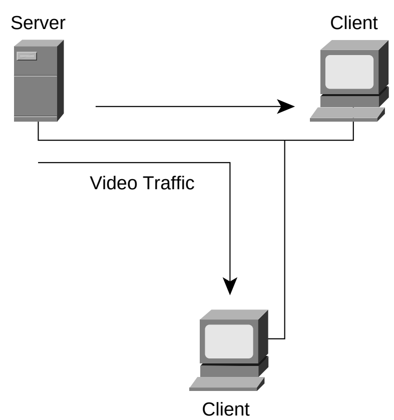

- Video server applications
- LAN TV

Both of these applications provide unidirectional video services. The video server delivers pre-compressed video steams to multiple clients. The LAN TV applications, on the other hand, deliver either stored video streams or real-time video from a camera source. Distance Learning, in which classes are video taped and then broadcasted over the LAN and WAN to remote employees is a popular example of a point-to-multipoint, unidirectional video application.

Network Bandwidth: The Reality

Perhaps the most crucial question, from a networking perspective, is how much bandwidth do all of these applications require. The answer is simple. It depends. It fact bandwidth requirements can range anywhere from 100Kbps all the way up to 70 or 80Mbytes per second for digital video.

The figure below graphs various network multimedia applications to the amount of bandwidth that they each require.

Figure 14.

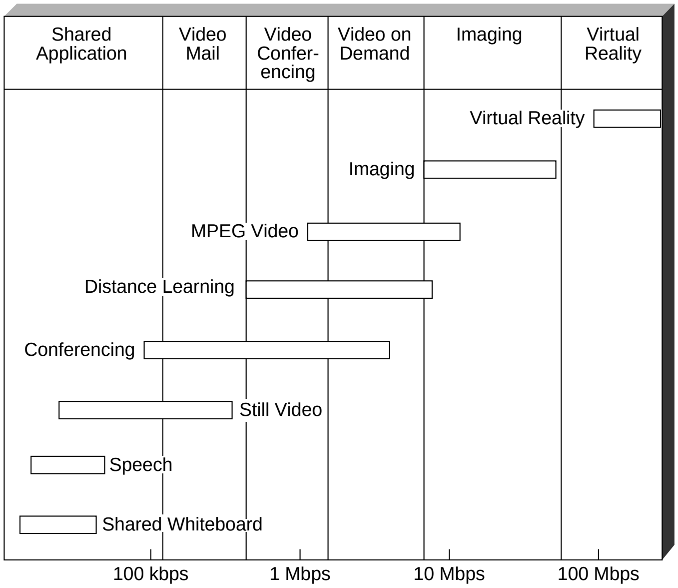

As Figure 14 indicates, the application selected has a direct impact on the amount of LAN or WAN bandwidth needed. Assuming that bandwidth is limited, the choice is either to select a lower quality video application that works within the available bandwidth, or consider modifying the network infrastructure to deliver more overall bandwidth. Section IV addresses the latter of the two choices, delivering more bandwidth to the Campus LAN and WAN.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Network Multimedia Applications

The table below shows some of the more popular network multimedia applications.

Table 5. Popular Network Multimedia Applications

Application	Type	Platform
Apple QuickTime Conferencing	Video Conf.	Mac
AT&T Vistium	Video Conf.	PC
CU-seeMe	Video Conf.	Mac/PC/UNIX
InPerson	Video Conf.	UNIX
Insoft Communique!	Video Conf.	PC/UNIX
Intel CNN at Work	LAN broadcast	PC
Intel ProShare	Video Conf.	PC
InVision	Video Conf.	PC
Novell Video for NetWare	Video server	NetWare
PictureTel	Video Conf.	PC
Starlight Starworks	Video server	UNIX/NetWare

Refer to the “Conclusion and References” section of this document for additional references on network multimedia applications.

IV. Bandwidth

Campus LAN Bandwidth

As the above discussion reveals, network multimedia applications range in their bandwidth consumption. In some Campus LAN environments, there is already adequate bandwidth available to run the desired network multimedia application. In most Campus LAN environments, this is not the case. In many situations this is a product of inefficient LAN design and segmentation, NOT a slow LAN medium. A considerable amount of bandwidth can be gained simply by re-segmenting, or micro-segmenting, the Campus LAN environment.

To illustrate the point, consider Campus A, Campus B and Campus C. Campus A has 500 users, on 5 separate 100-node shared Ethernet segments. Each of the five segments are connected via ethernet using a Cisco router. Based on 100 users per segment, the net bandwidth per user is roughly 100Kbps. Looking back at Figure 14, Campus A would at best only be able to handle an audio conferencing package.

Figure 15. Campus A Network Design: Shared Ethernet

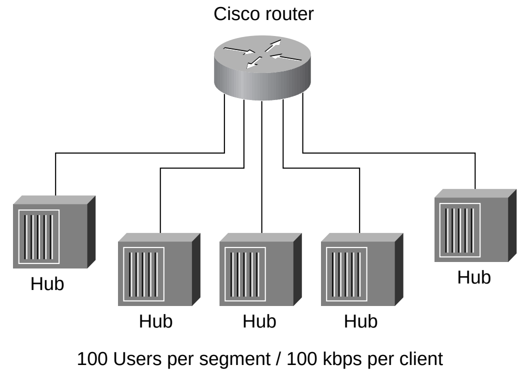

Campus B, on the other hand, uses a combination of share Ethernet hubs (repeaters) and Cisco Ethernet Switches to deliver substantially more bandwidth per client. In this design, 10 users are connected to a shared Ethernet hub. The hub is then connected to dedicated 10Mbps Ethernet switch ports on a Ethernet Switch. Each of the Switches are connected together over a routed Ethernet backbone. In this scenario, each hub gets 10Mbps which in turn yields roughly 1Mbps for each of the ten clients on the hub. Based on this network design, Campus B is capable of running medium to high quality video applications.

Figure 16. Campus B Network Design: Shared Ethernet/Switched Ethernet

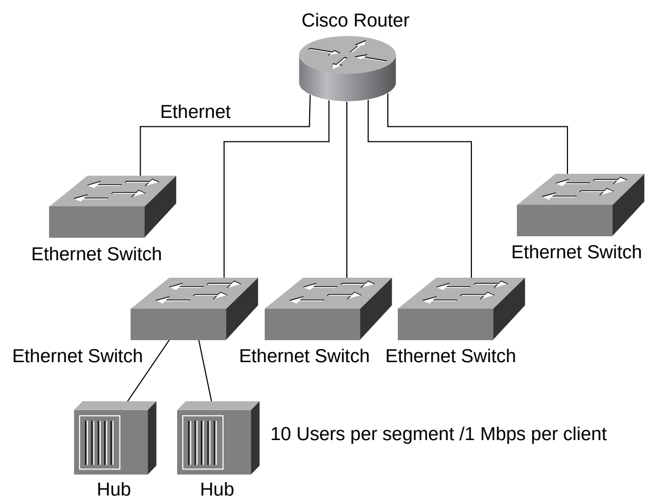

Campus C went one step further by eliminating the shared Ethernet hub altogether. Here, all clients have their own dedicated 10Mbps connection to the LAN by having a direct connection to an Ethernet switch port. Like Campus B, the switches are all interconnected over a routed Ethernet backbone. With 10Mbps per client, Campus C can easily support high-quality network multimedia applications.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 17. Campus C Network Design: Per Client Switched Ethernet

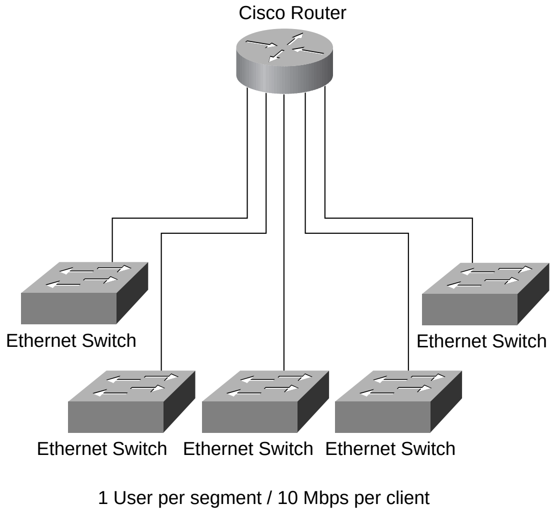

As the above example indicates, the first step in delivering more bandwidth is not removing the existing Ethernet or Token Ring infrastructure and moving to a 100Mbps technology. Rather, the proper first step is to micro-segment the LAN using switching. Remember that the majority of today's networked multimedia applications require less than 10Mbps for operation. Hence, Ethernet is still an acceptable to LAN medium; it is just that more of the 10Mbps needs to be delivered to each client than was the case before.

The following two figures shows how the effects of microsegmentation can increase change to end bandwidth which in turn allows running higher bandwidth network multimedia applications.

Figure 18. Typical Campus Usage Pattern

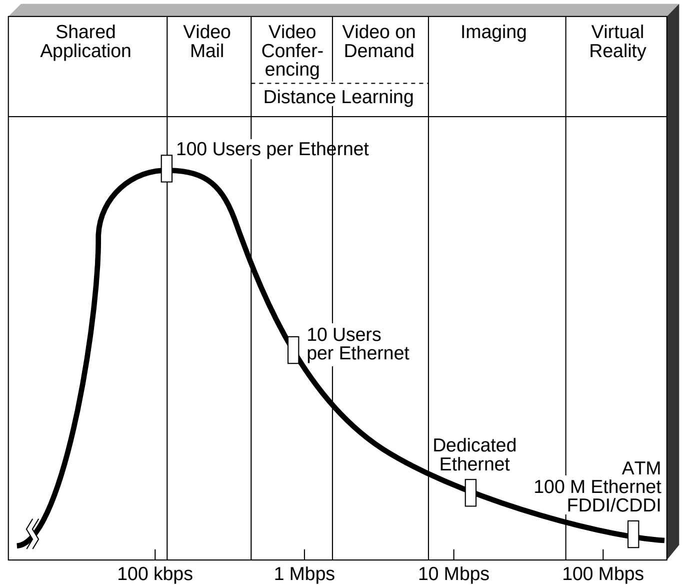

Figure 19. Segmented Campus LAN using Ethernet Switches

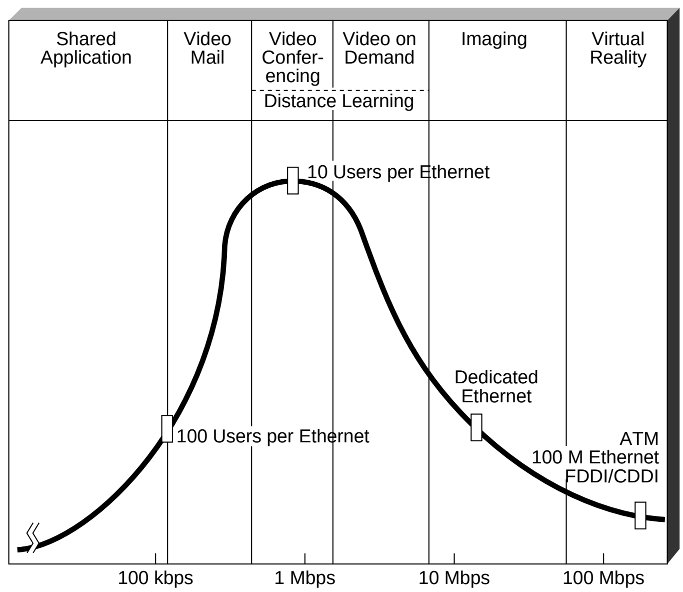

Cisco Ethernet Switching Products

Note: This section does not include information about Grand Junction products (Catalyst 1000 and 2000 series).

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Kalpana EtherSwitch Product Family

Cisco offers a range of products to help microsegment Campus LANs. At the entry level is the Kalpana family of layer 2 (Data Link) Ethernet switch products. Currently there are 3 products from Kalpana: EtherSwitch Pro16 (Catalyst 3000), Ether Switch EPS-2015RS and EtherSwitch EPS-2115M.

The Pro16 is the foundation for the ProStack, a high-performances, stackable switching platform. The Kalpana ProStack consists of the EtherSwitch Pro16 and the ProStack Matrix. The Pro16 supports 16 switched Ethernet ports with expansion slots for two high-speed modules. Each expansion slot can be populated with a mix of modules to suit user requirements. If high-speed connectivity is need, the Pro16 offers Fast Ethernet and ATM modules. From two to eight EtherSwitch Pro16 units can be stacked to achieve a system of up to 192 switched ports with a switching capacity of 4.8Gbps.

The EtherSwitch EPS-2015RS is a 15-port 10BaseT EtherSwitch with a slim, rack-and-stack design. This switch is designed specifically to complement the low-cost connectivity of stackable, multisegment hubs in departmental and workgroup 10BaseT networks.

Finally, there's the EPS-2115M which is a modular design that supports 10BaseT, 10Base2 and 10BaseFL. From an enterprise perspective, the Kalpana ProStack offers considerable flexibility, especially with its available high-speed uplinks.

Note that the Kalpana ProStack is available as either the Catalyst 3000 or the CiscoPro EtherSwitch Stack System.

Note that all Kalpana switches perform only Layer 2 switching operations. As a result, Layer 3 (Network Layer) information is ignored. Consequently, the network is more susceptible to broadcast storms because Layer 3 filtering is not possible. In small environments this is not a problem but in large campus environments this could have a crippling effect. Aside from limiting the number of switches per subnet, VLAN technology can also help to control broadcast and multicast traffic in switched networks. For more information regarding VLANs refer to the VLAN section that follows. And for an even more in-depth discussion refer to Harbrinder Kang's *Switched LAN Design Guide*.

Figure 20. Kalpana-based LAN Design

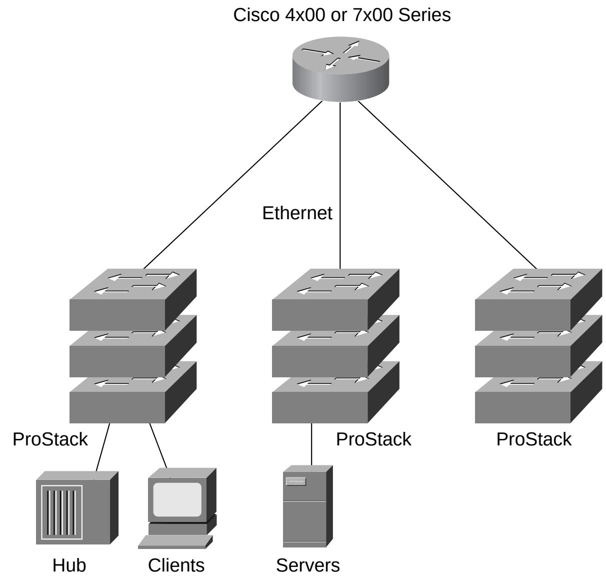

Catalyst Ethernet Switching Product Family

Cisco also offers the Catalyst family of Ethernet switch products: the Catalyst 1200 and the Catalyst 5000. The Catalyst 1200, not to be confused with the Catalyst 1700 from Grand Junction, is a multilayer switch for workgroup applications that can benefit from OSI layer 3 as well as Layer 2 capabilities. Layer 3 switching gives users the high performance of Ethernet switching combined with the flexibility, security and reliability of routing. The Catalyst 1200 offers 8 10BaseT or 10BaseFL ports and one expansion slot. The expansion can be populated with either 1 A/B CDDI interface or 1 A/B FDDI interface.

In addition to meeting a wide range of performance needs for Ethernet and FDDI, the Catalyst 1200 offers unique features that help managers monitor and control the growth and change of client/server workgroups. As part of this, the Catalyst offers embedded RMON functionality which offers protocol-decoding capabilities previously available only on full-function network-analyzers

The Catalyst 1200 also includes support for IGMP (Internet Group Management Protocol). IGMP is a standard-based protocol for group membership to multicast groups. Multicasting will be discussed later, but the point to remember here is that if IGMP is needed for a specific network multimedia application, it ideally should be supported on the switches that are in use. The Catalyst 1200 is one of only a few switches in the market that offers this support.

Figure 21 shows a typical Campus LAN design that uses Catalyst 1200s to deliver bandwidth to clients and servers.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 21. Catalyst 1200-based LAN Design

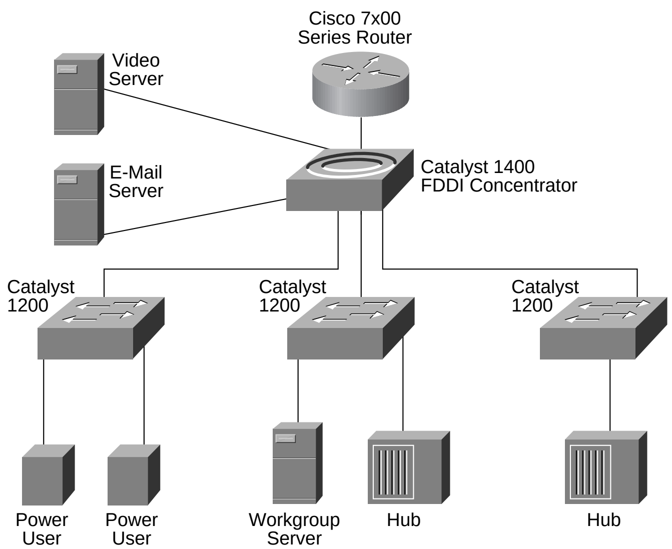

At the high-end is the Catalyst 5000, a modular switching platform that meets the needs of today's high performance, bandwidth-intensive network switching applications. The Catalyst 5000 offers five slots which can be populated with any combination of 10BaseT/10BaseFL modules, switched 100Mbps Fast Ethernet, FDDI or ATM modules.

The Catalyst 5000 is designed from the ground up to deliver high-performance both for client and server connections as well as for backbone connections. The Catalyst 5000's switching backplane operates at 1.2Gbps and provides nonblocking performance for all switched 10Mbps Ethernet interfaces.

Figure 22 shows the Catalyst 5000 in three different configuration scenarios.

Figure 22. Catalyst 5000-based LAN Design

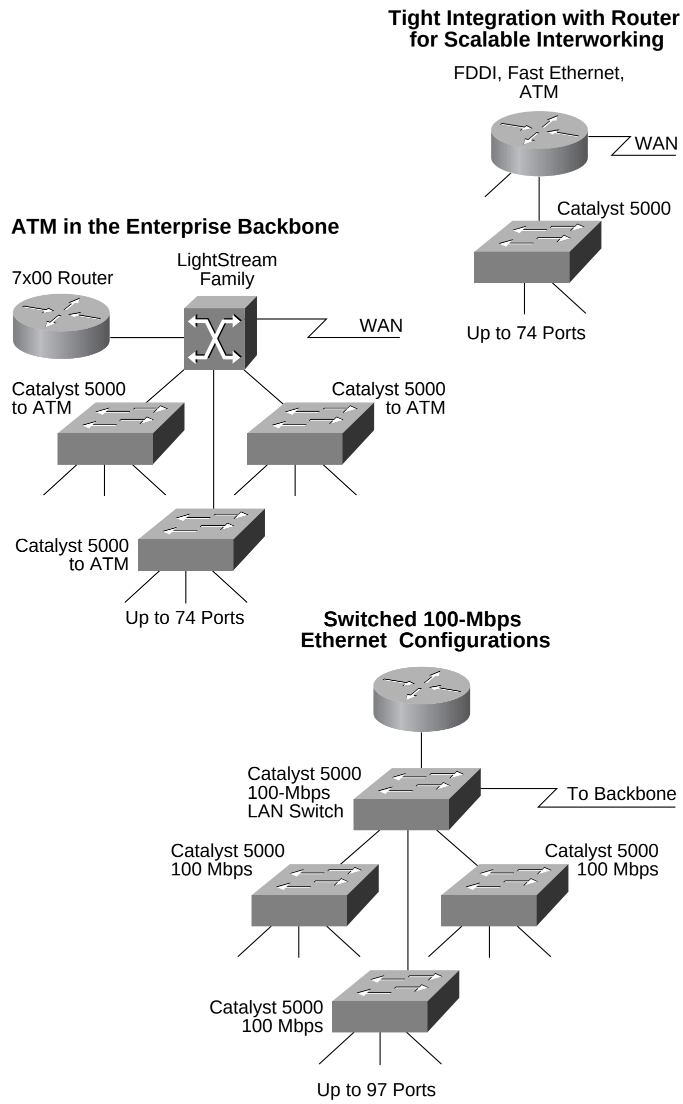

Token-Ring Switching: The Catalyst 1600

While Ethernet switching has gained the lion's share of attention in the market place, token-ring switching is also available. Similar to Ethernet switching which provides multiple 10Mbps Ethernet segments, token-ring switching provides multiple 4 or 16 Mbps token-rings.

The Catalyst 1600 is an eight port token-ring switch that was developed in conjunction with Madge Networks. Currently the Catalyst 1600 only supports token-ring, but support for FDDI and ATM modules are scheduled for September of 1995 and Q1 of 1996, respectively. As Figure 23 depicts, the Catalyst 1600 is instrumental in delivering greater bandwidth to both clients and servers.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 23. Catalyst 1600-based LAN Design

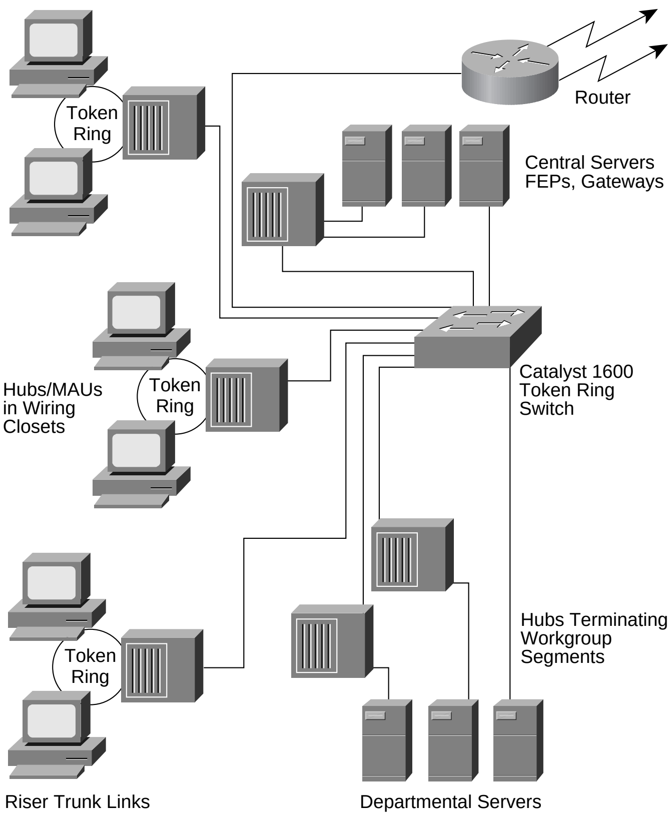

High-Speed LAN Technologies

As noted above, both the Kalpana and Catalyst families of switching products offer optional high-speed uplinks: Fast Ethernet, FDDI/CDDI and ATM.

Fast Ethernet

Fast Ethernet (IEEE 802.3u), delivers 100Mbps over either category 5 UTP or fiber optic cable. Like 10Mbps Ethernet, Fast Ethernet employs the CSMA/CD network access method. Perhaps the two best advantages of Fast Ethernet are that its relatively inexpensive (assuming category 5 UTP is present) and migration from traditional 10Mbps is simple.

Fast Ethernet can serve a variety of different environments:

- High-speed server connectivity
- High-speed client connectivity
- High-speed uplink for inter-switch communication
- High-speed backbone

High-speed server connectivity is a popular use for Fast Ethernet. In this scenario, servers (Novell NetWare 3.x, 4.x; Windows NT, UNIX servers) are on Fast Ethernet and transmit to clients connected either via Fast Ethernet or switched 10Mbps Ethernet. Fast Ethernet server connectivity works particularly well in video server environments where the server needs to deliver multiple video streams to its clients. Again the ability to take advantage of the high-speed connection is a product of the server's architecture and the OS (operating system) that is running. Novell NetWare, for example, can deliver substantial I/O caching which in turn makes NetWare a solid platform for video server applications.

Figure 24. Fast Ethernet Server Access

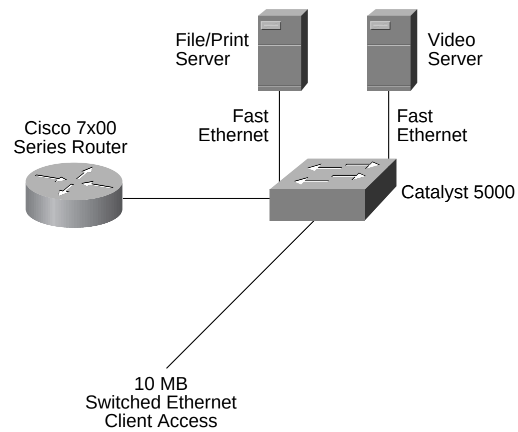

Using Fast Ethernet for high-speed client connectivity is also effective. Simply purchase a Fast Ethernet adapter for the client and connect it up to the Fast Ethernet switch. Today, Fast Ethernet adapters are available for PCs (EISA and PCI) and SPARCstations (S-bus) at very reasonable prices. And because installation is simply, Fast Ethernet provides a very straightforward migration path to 100Mbps.

Figure 25. Fast Ethernet to the Desktop

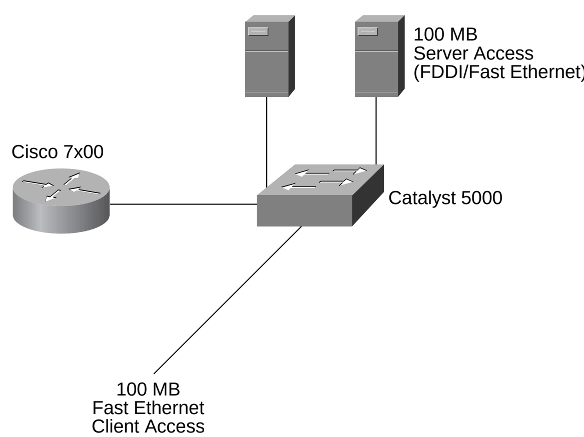

Fast Ethernet can also be used to interconnect Ethernet switch workgroups. In this scenario a group of Kalpana Pro16s, for example, are interconnected using Fast Ethernet. This is particularly useful in a micro-segmented environment where each client has a dedicated 10Mbps pipe. With a Fast Ethernet connection between switches, a client can communicate with

Cisco Confidential
Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

another client attached to a different switch without sacrificing bandwidth. If the inter-switch connection were only 10Mbps, the two clients would most likely sacrifice bandwidth assuming other access across the same inter-switch link. Fast Ethernet can also be used to interconnect Kalpana Pro16s to Catalyst 5000s or to interconnect Kalpana Pro16s and Catalyst 5000s to a Cisco 7000 series router equipped with a Fast Ethernet Interface Processor (FEIP). Again the benefit is to provide high-speed interswitch communication via a "switch of switches" or a routed network.

Figure 26. Fast Ethernet Inter-Switch Connectivity

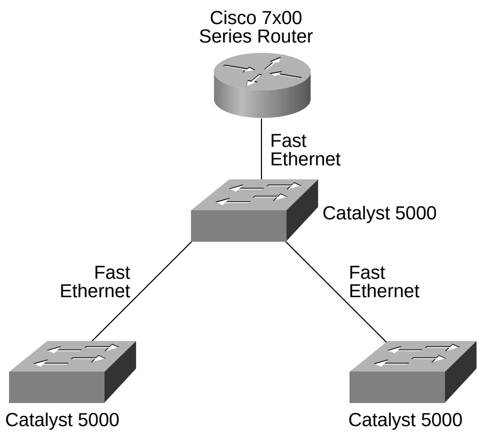

Finally, using fiber, Fast Ethernet can be used as a backbone technology to interconnect various switched segments in a Campus environment. This scenario is similar to the previous scenario, except the distances between switches and routers are greater here. This in turn requires that fiber be implemented because category 5 UTP limits Fast Ethernet connections to 100 meters. With fiber, however, Fast Ethernet can deliver connections up to 2 kilometers. In most cases, Fast Ethernet will not be used as a core backbone technology, giving way to FDDI and ATM. Simply put, FDDI and ATM offer advanced features that make them more viable for backbone implementations.

Figure 27. Fiber Fast Ethernet Backbone Design

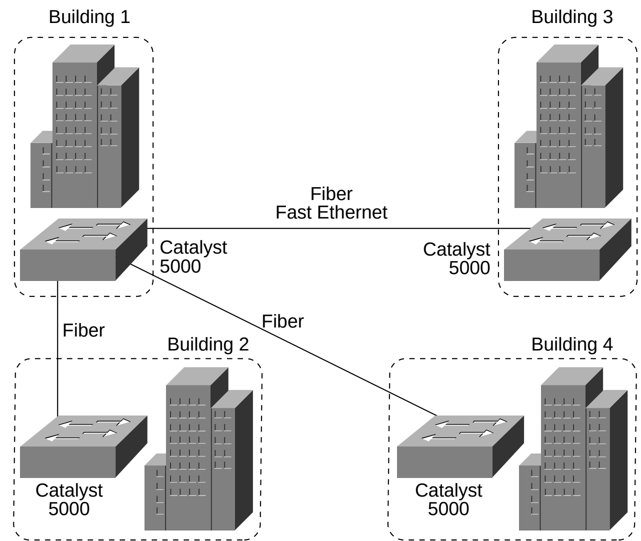

FDDI/CDDI

Both the Catalyst 1200 and the Catalyst 5000 offer high-speed uplinks using FDDI. Like Fast Ethernet this allows for a variety of different network design scenarios:

- High-speed client connectivity
- High-speed server connectivity
- High-speed uplink for inter-switch communication or to "switch of switches"
- High-speed backbone

Like with Fast Ethernet, FDDI can be used to deliver high-speed client connectivity. But in most cases, FDDI will be used for server connections and backbone connections. Especially in video server environments where multiple video streams are being sent out to video clients, FDDI/CDDI can deliver fast and reliable connections.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 28. FDDI Client and Server Connectivity

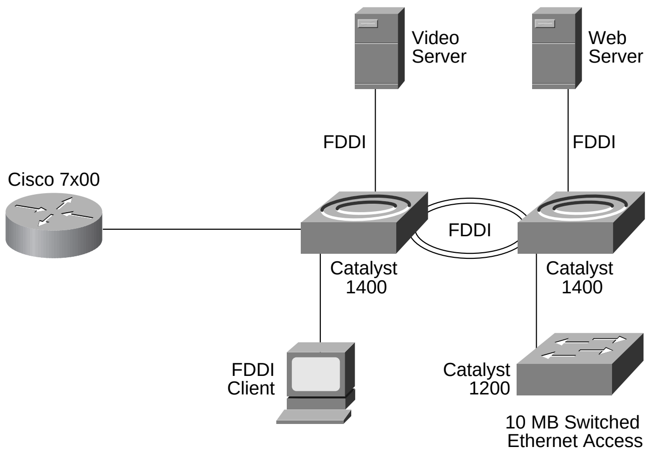

Besides delivering high bandwidth, FDDI/CDDI can also deliver greater redundancy especially when compared to Fast Ethernet. With FDDI or CDDI a server can be "dual-homed" to FDDI/CDDI concentrators. Dual homing provides a server access to two FDDI/CDDI rings. Under normal circumstances, the server will use only one ring. If for some reason, however, the primary ring fails, the server can fall back to the secondary ring, maintaining connectivity with no downtime. Dual homing a server requires that the server FDDI/CDDI adapter be a Dual Attached Station (DAS) adapter. FDDI/CDDI adapters are available as either SAS (Single Attached Station) or DAS. SAS adapters only offer one physical connection while DAS adapters provide two physical. In order to dual home, two connections are needed.

Figure 29. FDDI Dual-Homed Design

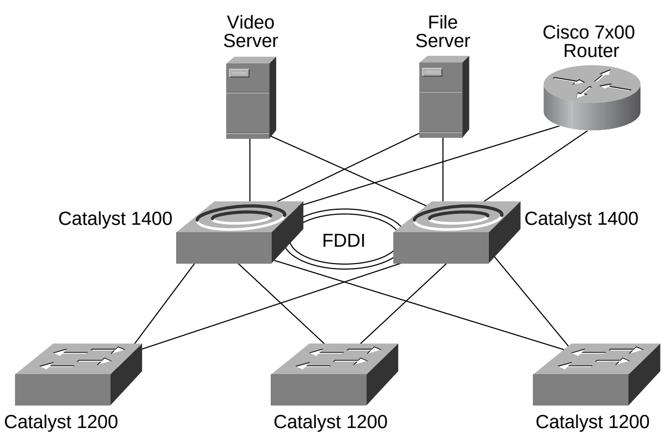

FDDI/CDDI can also be implemented to interconnect Ethernet switches. A group of Catalyst 1200s and of 5000s can be interconnected using FDDI concentrators (Cisco Workgroup Concentrator 1100 or Workgroup Concentrator 1400). Like with Fast Ethernet, the benefit here is that clients attached to different Ethernet switch workgroups can gain high-speed inter-communication. This in turn could permit an Ethernet switch client in one workgroup to access a video server or initiate a video conferencing session with a resource connected to a remote workgroup. Here again, dual homing can be implemented if desired. An FDDI-equipped switch can be dual homed to two separate concentrators providing greater redundancy and fault tolerance.

Figure 30. FDDI for Inter-switch Connectivity

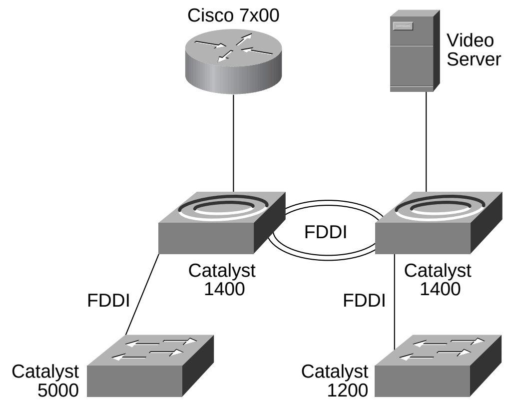

As in the above diagram, FDDI is particularly attractive as a backbone technology for three primary reasons:

1 Distance capabilities

2 Fault tolerance and redundancy

3 Security

There are two types of FDDI, multimode fiber and single mode fiber. With multimode fiber, a FDDI connection can span 2km. Single mode fiber, on the other hand, can span 10km. This in turn provides tremendous flexibility for interconnecting LAN segments in a Campus environment.

FDDI's inherent fault tolerance and its ability to support designs such as dual homing also makes the technology attractive in backbone environments. And from a security perspective, FDDI's optical transmission makes it more difficult for hackers to tap into, compared to traditional copper transmission.

ATM

Note: This section does not include information about the LightStream 1010.

Over the past few years, ATM has gained much attention as the next generation LAN and WAN technology. Much of the excitement in ATM centers around the fact that ATM delivers an entirely switched-based fabric and offers high-speed connectivity (100Mbps TAXI, 155Mbps OC-3 and in the future 622Mbps OC-12). Besides the raw bandwidth that ATM provides, the technology also offers extensive support for transporting video, voice and data through ATM-based Quality

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

of Service (QoS) parameters. This sections serves to provide an overview of Cisco's ATM-based solutions and design scenarios. ATM QoS will be discussed later in this document.

From a design and signaling standpoint, there are two key concepts to understand: UNI and NNI. The UNI, or User to Network Interface, specifies the connection between ATM end-systems (hosts, routers, etc.) and an ATM switch. NNI, or Network to Network Interface, specifies the connection between two ATM switches. Different cell formats are defined across UNI and NNI connections. Figure 31 below is a diagram depicting UNI and NNI boundaries in an ATM-based network.

Figure 31. ATM UNI and NNI Boundaries

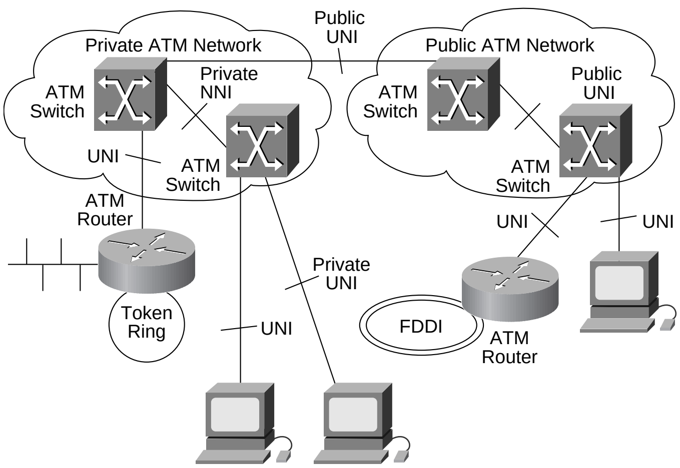

Because different signaling standards are used for different parts of an ATM network, and because not all signaling standards are available on all ATM equipment, it is important to know the equipment that is being used. Today, not all signaling standards are supported on all ATM products.

Currently Cisco offers six ATM-based products:

- LightStream 100
- LightStream 2020
- AIP for Cisco 7x00 series
- ATM NPM for Cisco 4000 series
- ATM Line Card for Catalyst 5000
- ATM S-bus/PCI LAN adapters

The LightStream 100 is an 16-slot ATM workgroup switch with support for the interfaces shown in the following table.

Table 6. LightStream 100 Interface Support

Interface	Line Speed	Medium	Connector Type
STS3c/STM1	155 Mbps	Multimode fiber	SC
TAXI 4B/5B	100 Mbps	Multimode fiber	MIC (FDDI style)
STS3c/STM1	155 Mbps	Single-mode fiber	SC
STS3c/STM1	155 Mbps	UTP-5	RJ-45
DS3	45 Mbps	Coaxial cable	BNC
E3	34 Mbps	Coaxial cable	BNC

The LightStream 100 supports signaling protocols that conform to the ATM Forum UNI version 3.0 specification. Future releases will support the pending ATM Forum UNI version 3.1 signaling protocol based on ITU-T Recommendations Q.2931 and Q.2110. The signaling will support point-to-point connection setup using any of the address formats defined by the ATM Forum, including E.164 or NSAP-encoded ATM private network addresses. A built-in segmentation and reassembly (SAR) function in the switch allows it to support ATM signaling and network management functions. Either AAL5 or AAL 3/4 can be used for carrying signaling requests.

In addition to supporting UNI signaling, the LightStream supports NNI functionality, enabling signaling requests to be routed in a multiswitch network. In the first release, the switch will support a prefix-based static routing protocol. As the P-NNI standards are developed, the Cisco IOS will be enhanced to support them. Because of built-in signaling support, the switch does not require a separate connection management system—thus lowering system costs and enhancing overall reliability.

While the LightStream 100 serves ATM workgroup environments, the LightStream 2020, on the other hand, is designed from the ground up to serve the enterprise ATM environments. The LightStream 2020 is a 12-slot chassis which provides the following features based on release 2.0 of its software:

- 2GB of switching capacity
- Support for trunks at digital rates ranging from 128 Kbps to 155 Mbps
- Support for user ports ranging in speed between 56 Kbps and 155 Mbps

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

- Sophisticated congestion avoidance and multiple classes of service
- High-throughput Frame Relay DCE switching
- LAN bridging
- ATM UNI service
- Support for frame-based proprietary protocols
- SNMP management

Cisco also offers ATM interfaces for the 4000 and 7x00 series routers and for the Catalyst 5000. The Catalyst 5000 ATM module (ATM LAN Emulation module) supports an OC3 interface (155 Mbps) with a wide range of media options—single-mode fiber, multimode fiber, and UTP in a future release. Currently the module supports LEC functions only (refer the LANE discussion later in this document for an explanation of these terms).

The Catalyst 5000, in conjunction with an ATM-equipped 4000 or 7000 Series router, provides users with a complete LAN emulation solution independent of external hardware or software. The Catalyst 5000 ATM line card will also interoperate with any third-party vendor's LEC/LES/BUS implementation that supports ATM Forum-defined LAN emulation standards. Lastly, the ATM module is fully compatible with ATM Forum-defined standards, including the 3.0 UNI specification, Q.2931 signaling protocols, and the LAN emulation standards. Future release will provide future support for UNI 3.1.

ATM cards for the 4000 and 7000 Series are also available. Both adapters support OC3, DS3, and E3 connectivity and UNI 3.0 signaling. Additionally, these interfaces support RFC 1577 (classical IP and ARP over ATM) and LAN Emulation client (LEC), server (LES), configuration server (LECS) and Broadcast Unknown Server (BUS) functionality. The 4000 and 7000 series ATM interfaces are crucial components when uniting routing and ATM environments.

And for ATM-attached clients and servers, Cisco resells Zeitnet's ATM LAN adapters. Currently, Cisco offers an S-bus adapter, and a PCI adapter will become available shortly. The adapters run at 155 Mbps over either fiber or UTP.

As Figure 32 illustrates, a variety of different design scenario are possible using Cisco ATM equipment.

Figure 32. Enterprise ATM Network Design

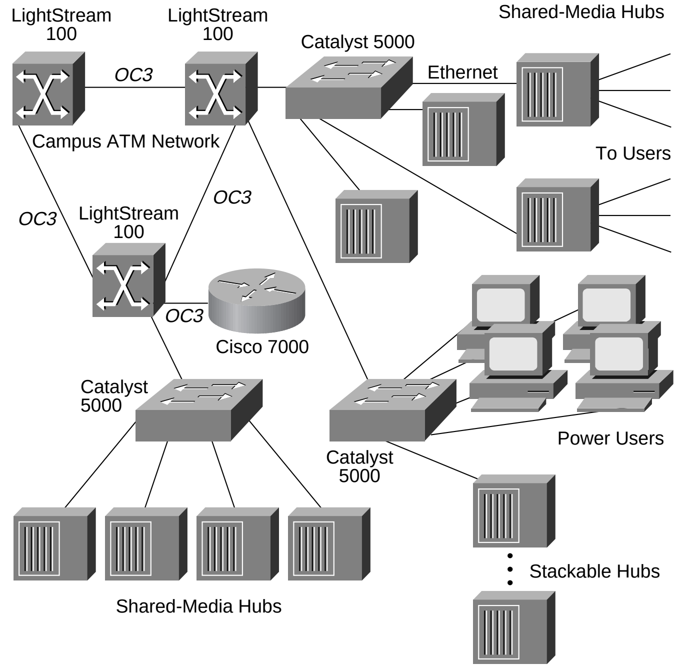

As for running network multimedia applications, ATM offers considerable flexibility from a bandwidth perspective. The point to remember, though, is that ATM is not a prerequisite to running network multimedia applications. Yes, there are advantages to running over ATM, such as ATM's inherent Quality of Service (QoS) capabilities. Nevertheless, today's existing LAN technologies can also support many network multimedia applications.

VLANs

Virtual LANs, VLANs for short, is another budding technology that can aid in the deployment of network multimedia applications. At the core of the VLAN is Ethernet and/or Token-Ring switching and ATM. VLAN technology adds to the aforementioned switching fabrics by allowing networks to be segmented into logically defined virtual workgroups. With VLANs, network managers can group switch ports and attached users into logically defined communities of interest. These groupings can be co-workers within the same department, a cross-functional product team, or diverse users sharing the same network application or software (such as a LAN Broadcast application). Grouping these ports and users into communities of interest, referred to as VLAN organizations, can be accomplished within a single switch, or more powerfully, between connected switches in a campus environment.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 33. Logical Networking Using VLANs

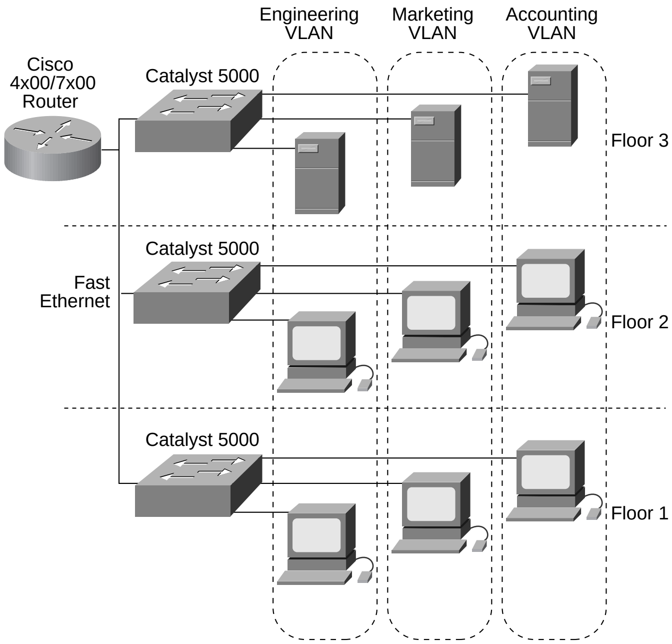

New client/server applications have driven the need for greater bandwidth in traditional shared-media environments and LAN switching has been viewed as the solution. LAN switches, however, dissolve previously well-defined workgroup/department boundaries because they behave like traditional “transparent” bridges, and all traffic flows through them. A network built and designed with LAN switches appears as a “flat” network topology; that is, a single broadcast domain. Consequently, in this environment, the networks are liable to suffer the problems inherent in “flat,” or bridged network.

In a “flat” network, all broadcast packets generated by any node in the network are sent to and received by all other nodes in the network. This may also be true for multicast packets. The ambient level of broadcasts generated by the higher-layer protocols in the network—known as “broadcast radiation”—will typically restrict the total number of nodes that the network can support. In addition bridging requires that all packets with unknown destination addresses be flooded, which can also undermine switch scalability.

VLANs have been designed to control broadcast traffic and to allow for more scalable switched network designs. A VLAN consists of a single bridged domain and solves the scalability problems of a large “flat” network by breaking a single bridged domain into several smaller bridged domains commonly referred to as VLANs.

Figures 34 and 35 show two different VLAN scenarios. The first figure shows intraswitch VLANs, 2 VLANs are created with one switch. The next figure shows both intraswitch VLANs and interswitch VLANs. Two different VLANs exist on one switch and both VLANs span the campus environment.

Figure 34. Intra-Switch VLANs

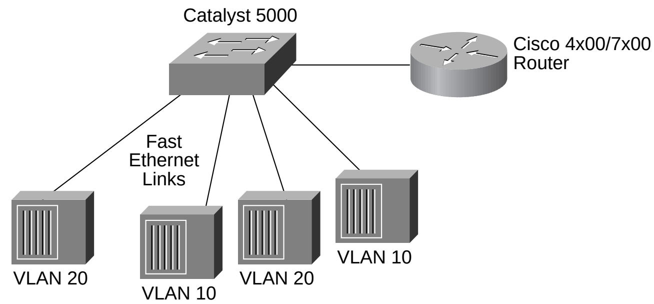

Figure 35. Inter-Switch VLANs

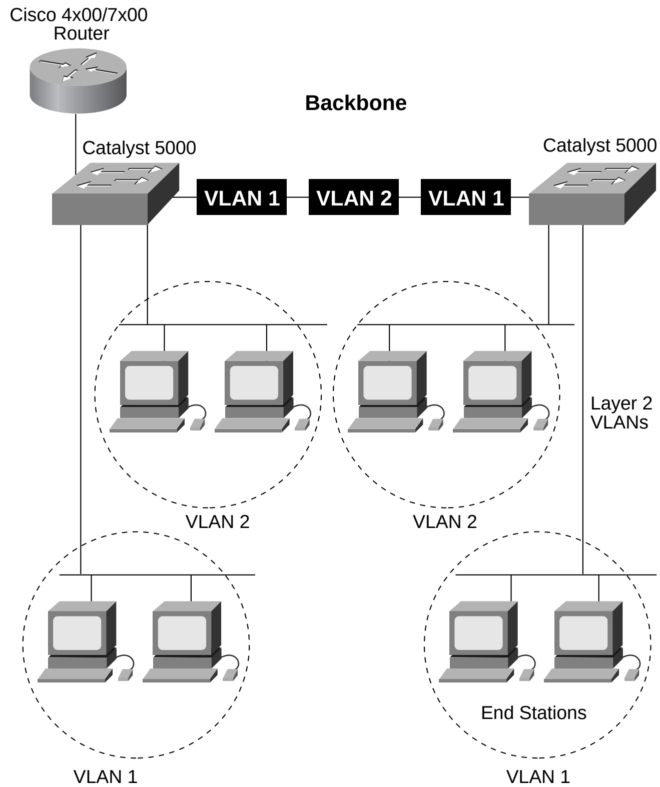

VLAN Support and Interoperability

Building campus-wide VLANs can be accomplished using fast Ethernet, FDDI, or ATM and their associated VLAN technologies: ISL (Inter-Switch Link), 802.10, and LAN emulation.

Currently VLAN support is available for the Kalpana ProStack 16, the Catalyst 1200, the Catalyst 1600, the Catalyst 5000, the LightStream 100 and Cisco routers. The table below shows available VLAN standards by Cisco platform.

Cisco Confidential
Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Table 7. Cisco Platform Available VLAN Standards

	ISL	802.10	LAN Emulation	Kalpana® (Proprietary)	LightStream™ (Proprietary)
Catalyst™ 1200	No	Yes	No	No	No
Catalyst 5000	Yes, FEIP	Yes, FDDI	Yes, ATM	No	No
ProStack	Q1’96	Future	ATM PVC (Q3’95) ATM LAN Emulation (Q1’96)	Yes	No
Routers	FEIP, Cisco IOS Release 11.1	All LANs, Cisco IOS Release 11.1	Yes, Cisco IOS Release 11.1	No	No
LightStream 100	No	No	Yes, ILMI	No	No
LightStream 2020	Future	Future	Future	No	Yes
Zietnet	N/A	N/A	Yes	N/A	N/A

Depending on which Cisco platforms are deployed, a wide-range of Campus LAN VLAN architectures can be deployed. For additional information on design VLAN-based networks, refer to *Designing Switched LANs* authored by Harbrinder Kang.

Routers in Switching Environments

Although VLANs will solve some of the broadcast and security issues in today’s traditional LAN environments, it is important to clarify that VLANs without routers will not scale to large campus environments. Routing is instrumental in the building of scalable VLANs and is the only way to impose hierarchy on the switched VLAN internetwork.

A *router* is a device that connects heterogeneous networks. Routing functions are performed per-protocol, at layer 3 of the OSI 7-layer model. Bridging or switching functions act on the OSI layer 2 source and destination addresses in the datagrams. A key point to note is that routers permit greater scalability through address summarization (or hierarchical addressing), because layer 3 addresses typically have structured (IP subnets). “Flat” addressing *requires* flooding, which limits scalability.

With routers, hosts using protocols with network layer addressing can immediately reduce the problem of responding to another host in a large network. If the destination has the same network number or ID, the host can respond to the locally attached network. If the destination network is unknown or cached in the host table as a “remote” network, the datagram can be sent to the router. The router then ensures that the datagram is sent to the remote network through the proper path.

The biggest benefit of a router is that destinations and the routes to them through the network no longer have to be discovered by every host on the network. The immediate effect is that the amount of broadcast and multicast traffic sent by a given host is reduced still further, and discovery and advertising traffic is NOT forwarded to other network segments. Instead the router caches advertised network services and destinations, forwarding datagrams appropriately when requested to do so by hosts on various segments.

Large router networks can effectively use redundant paths and have informational protocols that allow correct path selection, even in a dynamically changing network. These routing protocols eliminate the need for packet flooding (transmitting unknown destination packets out all ports) and network-wide broadcasts for discovery or advertisement, for the reasons stated previously; address summarization.

Below are two network designs a that employ both routers and switches.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Figure 36. Router/VLAN Integration Example 1

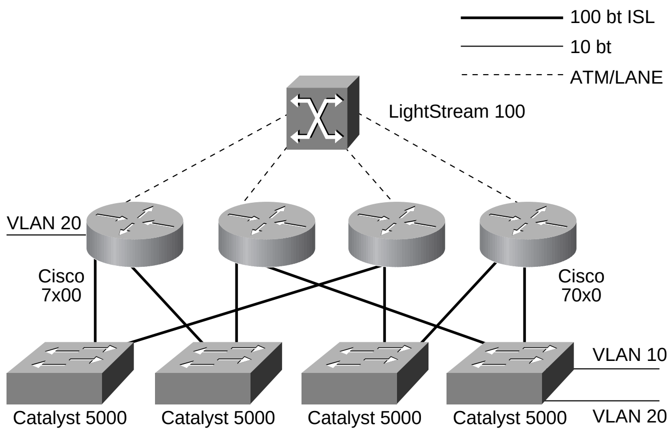

Figure 37. Router/VLAN Integration Example 2

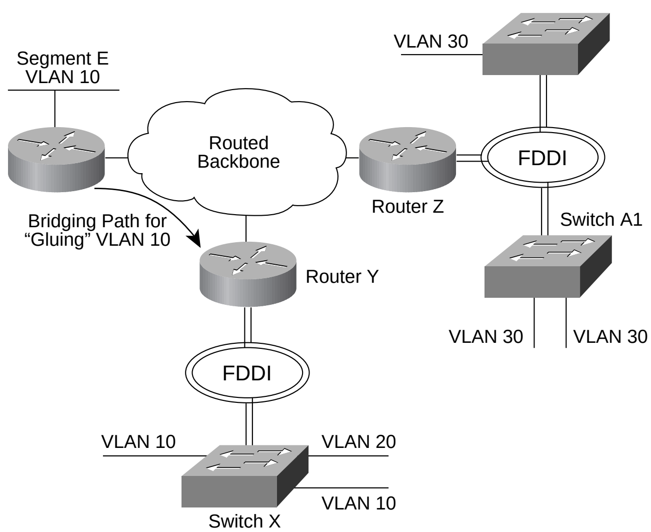

WAN Bandwidth Considerations

While there are many different ways of increasing LAN bandwidth, increasing WAN bandwidth is not as easy. In fact WAN bandwidth in many environments is a scarce resource. This is due mainly to one simple point: WAN bandwidth is more expensive. Consider that LAN technologies are measured in Megabits while WAN bandwidths are mostly measured in Kilobits. Even a T1 (1.544 Mbps) which is a good size in the WAN environment is 6 and a half times slower than Ethernet!

As for running multimedia applications across the WAN, well, it's a challenge, to say the least.

If additional bandwidth is needed in the WAN, look first at available circuit-switched technologies: switched-56, switched-T1, and ISDN in conjunction with the Cisco IOS DDR feature. With these services, charges are based on connect time. This means if they are used for multimedia services, charge will be based on the duration of the multimedia session. In those cases where the circuit switched service is used in conjunction with another connecting WAN service (switched or leased), the circuit-switched service can be configured as a backup service for both link failure and bandwidth needs.

The key issue to remember for circuit-switched services is that charges are incurred based on usage. This means that in order to curb costs, uptime must be controlled. Cisco IOS offers a variety of features to help control DDR links. Some of these features include:

- Snapshot routing
- IPX spoofing
- SPX spoofing
- NBP filtering
- IP, IPX, AppleTalk, VinesIP, DECnet access lists
- Policy-based routing

All of these features are instrumental in controlling WAN uptime across circuit-switched connections. For more information on these features refer to Cisco Connection Documentation CD-ROM. For ISDN specific design and implementation information, refer to the *ISDN Design and Implementation Guide* (J. Baher).

In addition to using circuit-switched WAN links to add bandwidth in the WAN, there are additional Cisco IOS Bandwidth-on-Demand (BOD) features that can further help the situation. In particular, ISDN b-channel load balancing and dialer rotary groups can help dynamically increase bandwidth between sites.

ISDN b-channel load balancing is accomplished using the **dialer-load threshold** command. This command allows additional b-channels to be brought up between two sites. In a ISDN BRI environment, two connected sites can take advantage of both 64Kbps b-channels, yield an aggregate bandwidth of 128Kbps. Using the dialer-load threshold command, a second b-channel is brought when the first channel exceeds a user-defined threshold. The threshold can be set to 0 so that the second b-channel is immediately brought up.

Note that Cisco IOS 11.0(3) and later supports PPP multilink (RFC 1717) for ISDN b-channel aggregation.

The following is a configuration example showing the dialer-load threshold command:

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

Commands in “[ ]” are available in Cisco IOS 11.0(3) and later.

```
!
hostname bertha
!
enable password ######
!
username bertha password 7 2394943E02B17
isdn switch-type basic-5ess
!
interface Ethernet0
ip address 171.68.158.49 255.255.255.248
!
interface BRI0
ip address 171.68.158.26 255.255.255.248
encapsulation PPP
dialer map ip 171.68.158.25 name althea speed 56 14085551111
dialer load-threshold 128 [inbound/outbound/either]:used to specify
    traffic direction
[ppp multilink]:enables ppp multilink
dialer-group 1
ppp authentication chap
!
ip route 131.108.0.0 255.255.0.0 171.68.158.25
ip route 171.69.0.0 255.255.0.0 171.68.158.25
access-list 101 permit ip any any
!
dialer-list 1 list 101
```

Another Cisco IOS feature that helps to dynamically add bandwidth to the WAN is the dialer rotary group feature. Like the dialer load-threshold feature which aggregates multiple b-channels, the dialer rotary group feature aggregates multiple circuit-switched connections. Designed to work with switched-56 and ISDN, the dialer rotary group feature can bring up multiple switched-56 or ISDN connections between sites. In other words multiple 56Kbps links can be brought up or multiple ISDN links can be brought up based on user defined parameters.

The following is an example of dialer rotary groups being used to interconnect to sites.

```
!
int bri 0
dialer rotary-group 1
int bri 2
dialer rotary-group 1
int dialer 1
ip address 3.3.3.3 255.255.255.0
dialer map ip 3.3.3.4 name icarus speed 56 4085551212
!
```

The Cisco IOS 11.0 release offers another solution, policy-based routing, for using circuit-switched WAN services in conjunction with leased line services. With policy-based routing, traffic can be routed over redundant WAN links based on traffic type (protocol, UDP port number, etc.). For instance, Policy-based routing can be used to route e-mail and FTP traffic over the serial link and route Intel ProShare traffic across the ISDN link.

The following is an illustration of Policy-based routing using a T1 interface for regular traffic and an ISDN interface for video conferencing traffic.

Figure 38. Policy-Based Routing

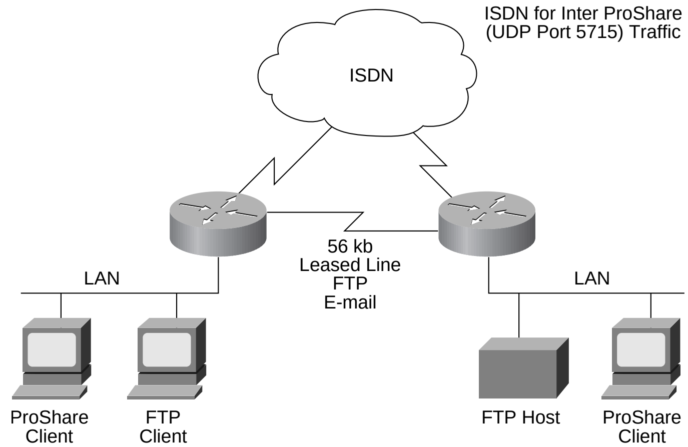

In this example, the multimedia traffic gets the required bandwidth from the circuit-switched service. Since the circuit-switched service is only up when the application is in use, WAN costs are contained. Traditional LAN traffic runs separately on the leased line an experiences uninterrupted service.

Until WAN bandwidth becomes affordable at any speed, delivering bandwidth to applications over the WAN will remain a difficult task. Wherever possible take advantage of circuit-switched technologies and Cisco IOS features like Policy-based routing, DDR, and BOD features.

Another point to consider is that not all network multimedia applications need to run over the WAN in real time. Naturally on-demand applications like conferencing will have to consume WAN bandwidth during the working day, but video server applications, for instance, do not. A typical video server environment might have multiple video servers deployed in various sights. During the day, users access their local video servers for training material or other video feeds. At night, when the WAN is for the most part idle, the video servers replicate information and receive updates of new video content. In this example, the network multimedia applications make use of unutilized WAN bandwidth while still not adding to the working day traffic load.

Cisco IOS Data Compression

Cisco IOS offers two different data compression algorithms for transmitting over WAN links: STAC and Predictor. While both algorithms are designed to “cram” more data onto WAN links, don’t think about using these algorithms for further compressing video and audio traffic. These two lossless algorithms cannot do what video and audio compression

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.

algorithms can do. An uncompressed video stream will not benefit nearly as much by going through STAC or Predictor as compared to going through MJPEG or Indeo compression. In most cases, these algorithms will increase the video or audio file size rather than decrease it.

Copyright © 1995 Cisco Systems, Inc. All Rights Reserved.