ROBOTEL

Application Note

Revision Date: 2007-05-04
Creation Date: 2006-01-05
Author: Gerry Sullivan
File Name: AppNote_Guidelines for Multimedia Servers & Networks_2006-02-27.doc

Guidelines for Multimedia Servers & Networks

INTRODUCTION

In language labs and multimedia classrooms, teachers often need to create a library of multimedia content that both teachers and students can access for different learning activities. In modern facilities, these media libraries comprise a set of computer files that are stored on a central media server.

Additionally, some language lab applications can be run from a central application server (rather than being installed on each student computer).

This application note examines the requirements for the media server, the application server, and the associated computer networks that link the servers with student and teacher workstations.

The main topics discussed are:

- Multimedia content
- The Media Server
- The Computer Network
- The Application Server

Most computer networks are not inherently optimized for the smooth delivery of multimedia content to a large number of users, so it is important that the topics addressed herein be properly understood and addressed.

MULTIMEDIA CONTENT

Media files contain some combination of audio and video information. When a media file is played using a media player (such as Microsoft's Windows Media Player or Apple's QuickTime), the user can listen to the audio content and/or view the video content.

Media files are most often created by *digitizing* analog signals from sources, such as a microphone, a video camera, an audiocassette tape, a videotape, etc.   A computer program – e.g. Windows Sound Recorder – is used to manage the digitizing process, which takes place in the computer's sound card or in an optional video-digitizing module.

Depending on the program and hardware used, the resultant digital file uses a specific data format. There are a wide variety of media file formats, and as computers get faster and information compression algorithms become more sophisticated, the list of new media file formats will continue to grow.

The Digitizing Process

When audio or video is digitized, we use a set of discrete digital samples to model the analog signal. Two of the main characteristics of the digitizing process are the *sample frequency* (how often a sample is taken) and the *sample size* (how much detail is captured for each sample).

A mathematician by the name of Nyquist discovered that the sample frequency used for digitizing must be at least 2 times the frequency of the program material contained in the analog signal. So, for example, if we wanted to digitize a high-fidelity audio signal containing program materials in the full range of human hearing – 20 Hz to 20,000 Hz – then our sample frequency must be at least 40,000 Hz. In fact, the audio CDs that you purchase from the local music store, use a sample frequency of 44,100 Hz (or 44.1 KHz).

Sample size, measured in bits, determines how precisely we want to capture the signal level. Again, using our audio example, this precision is called the dynamic range. Audio CDs use a 16-bit sample size, but other applications can use less (8 bits for telephony) or more (24 bits for recording studios).

The other point to consider is whether the analog audio signal is monophonic (1-channel) or stereo (2-channel). Some modern surround sound applications actually use up to 7 discrete audio channels.

But let's keep things simple. Imagine that we want to digitize an audio signal similar to those used on an audio CD. We will use a sample frequency of 44.1 KHz, and a sample size of 16 bits/channel for two stereo channels. If we do the math, we will discover that our resultant digital audio signal contains a significant amount of data. In fact, it contains about 10 Mbytes of data for every minute of audio. (We might also say that the signal bandwidth is 10 Mbytes/minute – or about 170 Kbytes/sec.)

Why is this important? The characteristics of digital media files are important, because our media server has to store this data, and both our media server and our computer network must be able to deliver this volume of data to system users continuously without interruption.

Signal Compression

To reduce the bandwidth and storage space required by media files, a number of clever folks, (like the Motion Pictures Experts Group, or MPEG for short), have devised ways of compressing raw media file data.

Some compression algorithms are *lossless*, meaning that one can fully recover the original signal without any loss, while other compression algorithms are said to be *lossy*, where some characteristics of the original signal are compromised.

For example, one popular compressed audio format is MP3, (an abbreviation for MPEG Layer 3). If audio is digitized in an MP3 format rather than a WAV format, one can achieve an effective compression of more than 10:1, without noticeable loss in signal quality.

It is important to understand what specific file formats you are using, and identify the size and bandwidth characteristics of these files.

Symposium & LogoLAB Digital File Characteristics

Lesson Materials

With the Symposium family of products, A/V capture files are digitized with a WMA format, (Windows Media Audio), or WMV format, (Windows Media Video). LogoLAB uses WAV for audio and WMV for video. The typical digital signal characteristics are shown in the table below:

File Type	File Format	File Bandwidth	Approximate File Size (Mbytes)		
			1 minute	10 minutes	100 minutes
Symposium Lesson Audio	WMA	128 Kbps	0.94	9.4	94
Symposium Lesson Video	WMV	450 Kbps	3.3	33	333
LogoLAB Lesson Audio (maximum)	WAV	1411 Kbps	10.3	103	1030

Student Responses

Symposium uses an MP3 format for open recording and a WAV format for other activities. LogoLAB also uses a WAV format. This is true for local student recordings and for session recordings saved by the teacher.

File Type	File Format	File Bandwidth	Approximate File Size (Mbytes)		
			1 minute	10 minutes	100 minutes
Student Audio (open)	MP3	128 Kbps	0.94	9.4	94
Student Audio (simultaneous or AAC)	WAV	1411 Kbps	10.3	103	1030

Other Media Types

As Symposium enables teachers to import a wide variety of digital file types, we must also consider the characteristics of other file types. Below we consider some common file types.

File Type	File Format	File Bandwidth	Approximate File Size (Mbytes)		
			1 minute	10 minutes	100 minutes
MPEG 1 Video	MPG	1700 Kbps	12.5	125	1250
MPEG 2 Video	MPG	3000 Kbps	22.1	221	2210

Storage Implications

To appreciate how much storage space we might need for a combination of lesson materials and student responses, let us consider a “typical” language lab comprising 32 student stations.

Let us further postulate that classes in this lab each take 90 minutes, and we run 4 classes per day, 2 in the morning, and 2 in the afternoon. Finally, let us assume that in each class, the teacher has each student save 5 minutes of audio responses.

Given these assumptions, we can then look at how much storage space is required for student responses over a semester of 12 weeks.

As it turns out, we will have:

12 weeks x 5 days x 4 classes x 32 students x 5 minutes x 10.3 Mbytes/minute = 384 Gbytes

Additionally, let us consider how much space is required for lesson materials. We shall assume that we have 10 teachers, and that each teacher has accumulated 20 hours of audio program materials and 10 hours of video program materials.

Given these assumptions, we can look at how much storage space is required for lesson materials:

8 teachers x (20 hours x 0.94 Mbytes/minute + 10 hours x 3.3 Mbytes/minute) x 60 minutes = 24 Gbytes

When we combine our lesson materials and one semester worth of student responses, our total storage requirements are about 408 Gbytes. More storage will be required if teachers use additional content and/or more than one semester’s worth of student responses are retained.

THE MEDIA SERVER

A *server* is actually a type of software service that runs on a computer, but most people tend to think of the computer that runs the particular service as *the server*. We shall assume that our media server process runs on a single computer called the Media Server.

We shall also assume that our Media Server is used strictly for recording, storing, and delivering media files, and that there are no other software application running on this computer.

Media Server Storage

In the previous section, we learned that our Media Server will need to have a hard drive at least large enough to store 408 Gbytes worth of lesson materials and student responses. Given the relatively low cost of hard drive storage, we should consider a single 160 Gbyte hard drive to be our minimum starting point for Media Server storage, but obviously to support all recordings for one term, we need as a minimum 2 disks of 250 Gbytes each.

Media Server Bandwidth

But storage is only one factor that we need to consider. The other significant issue is *bandwidth*. In our example class of 32 students, the server must be capable of concurrently delivering video files to all students – and we must consider the most demanding of file types. We shall assume that our worst case content is MPEG 1. The calculation of required minimum server bandwidth is:

32 students x 1700 Kbps = 53 Mbps

Our Media Server must be able to support a sustained throughput of 53 Megabits/second.

Unfortunately, computers do not usually advertise a *bandwidth* specification. Instead, we must look inside the server and consider those factors that act as bottlenecks to throughput. With modern computers, there are two major potential bottlenecks; network connections and hard drives.

Network Connections

In the previous section, we noted that for our example lab, the media server must support 53 Mbps. Fortunately, this is comfortably within the support envelope of a standard 100 Mbps Fast Ethernet network connection.

But, imagine if our lab doubled in size to 64 student stations, and we then needed to support 106 Mbps. This could not be done over Fast Ethernet. Instead, we would need to consider migrating to multiple network connections or using a faster network connection like Gigabit Ethernet, (1000 Mbps).

Hard Drives

A *hot* hard drive, like a 10,000-RPM Western Digital *Raptor*, supports a sustained throughput of 576 Mbps (72 MBps). This is more than enough for our lab example. Even most 7200-RPM drives can readily handle our required throughput of 53 Mbps.

However, we can imagine that if our server is shared with multiple labs, it is certainly possible for the required throughput to exceed the capabilities of a single hard drive.

RAID

RAID is an acronym for Redundant Array of Independent Disks. The idea behind RAID is that multiple hard drives can be combined in a way that their bandwidth is aggregated in a linear fashion. This means that (under ideal conditions) 2 hard drives have twice the bandwidth of a single drive.

That being said, not all RAID arrays are designed to optimize bandwidth. Some arrays are configured to optimize data redundancy (and therefore fault tolerance). Usually, some combination of improved bandwidth and fault tolerance is desired when migrating to a multi-disk RAID array. There are six traditional levels of RAID arrays, numbered 0 through 5:

RAID Array Type	Description
RAID 0	A RAID 0 array *stripes* data across each hard drive in the array. This provides improved bandwidth but without any fault tolerance.
RAID 1	A RAID 1 array mirrors data on two hard drives using a byte-striping approach. This provides fault tolerance and improved reading bandwidth, but does not improve writing bandwidth.
RAID 2	RAID 2 is theoretically designed for error correction, but it uses an unusual bit-striping approach that is rarely implemented in practice.
RAID 3	RAID 3 uses a more practical byte-striping approach, but because one disk in the array must hold only parity information, the use of RAID 3 is not common.
RAID 4	RAID 4 uses a block-striping approach and also uses one dedicated disk for parity information. This approach limits writing bandwidth
RAID 5	RAID 5, which uses a byte-striping approach, distributes parity information over all drives in the array and is the most popular choice for RAID implementations. It provides an excellent balance of bandwidth and fault tolerance.

In environments having a large amount of data and/or a large number of users, most media servers will include a RAID 5 array to support improved throughput and fault tolerance.

OPERATING SYSTEM

With conventional operating systems like Windows 2000 Professional or Windows XP Professional, the maximum number of concurrent users that can share a media file is dictated by Microsoft to be 10. To permit the sharing of media files with more than 10 users, a different operating system must be used.

In a Microsoft environment, the media server will typically run the Windows 2003 Server operating system. It is also possible to run Linux (with Samba support) on the media server – which after all is really just a file server.

With Robotel's language lab products, (Symposium, LogoLAB, and IntraLAB), it is typically only the Document folder (which stores multimedia documents for teachers) and the Student folder (which stores student responses) that are shared from the media server. Each product includes a set of configurable paths for these folders.

With Symposium and IntraLAB, the main application and its associated database can be run on the control computer under a conventional operating system like Windows XP Pro.

With LogoLAB, the main application must also be run on a server capable of supporting access from all concurrent users. In smaller environments, (e.g. 80 or fewer users), both the LogoLAB application and the LogoLAB media files can be run on a common server. In larger systems, the LogoLAB media files are run on a Media Server, and the LogoLAB application is run on a separate Application Server.

THE APPLICATION SERVER

LogoLAB uses a true application server model. LogoLAB is installed only on an application server and not on client stations. When several users are running LogoLAB, each instance of the program runs directly on the application server.

Normally, the Application Server will use a very specific operating system, Microsoft Windows 2003 Server, and the server must be licensed to support the required number of concurrent client accesses. It is however possible to run the Application Server under another operating system (e.g. Novell or Linux) if the associated Symposium LogoLAB database module is installed on a separate Windows computer, (WinXP, Win2000, Win2003 Server).

The same physical computer can be used as both an application server and as a media server, however for larger labs, it is strongly recommended that these two functions run on separate machines to ensure that computer resources (such as hard drives) are not overwhelmed by the combined burden of file serving and application serving.

THE COMPUTER NETWORK

The computer network is the backbone of a communications system that enables users to communicate with other users within the same facility (Local Area Network or LAN) and in remote locations (Wide Area Network or WAN). Most of us are also familiar with the WAN called the *Internet*, which permits us to access a wealth of data on the World Wide Web, and communicate with colleagues around the globe using programs like e-mail, instant messaging, and VoIP (Voice over Internet Protocol) telephony.

Computer networks were originally conceived to handle the transfer of data from one computer to another computer, and so adopted an infrastructure that was not too concerned with minor delays in moving data from point A to point B. With applications like telephony – which presume real time audio transmission – delays become a problem.

It is therefore critical, that within our language lab, the computer network be configured with enough useable bandwidth to conservatively handle all of the real-time (actually called *isochronous*) multimedia traffic that we are projecting without noticeable delay.

It is also important that the computer network within the language lab be somewhat isolated from the rest of the facility, so that the applications that others outside of the lab may be running do not materially affect our internal traffic – and vice versa.

Within the lab, one of the main tasks of the network is to support the delivery of content from the media server to the client stations (when playing media files) and from the clients to the media server when recording student responses.

As we had determined when looking at Media Server requirements, the minimal network configuration that we require must support at least a 100 Mbps connection to the Media Server and 10 Mbps connections to the client stations. This is shown visually below:

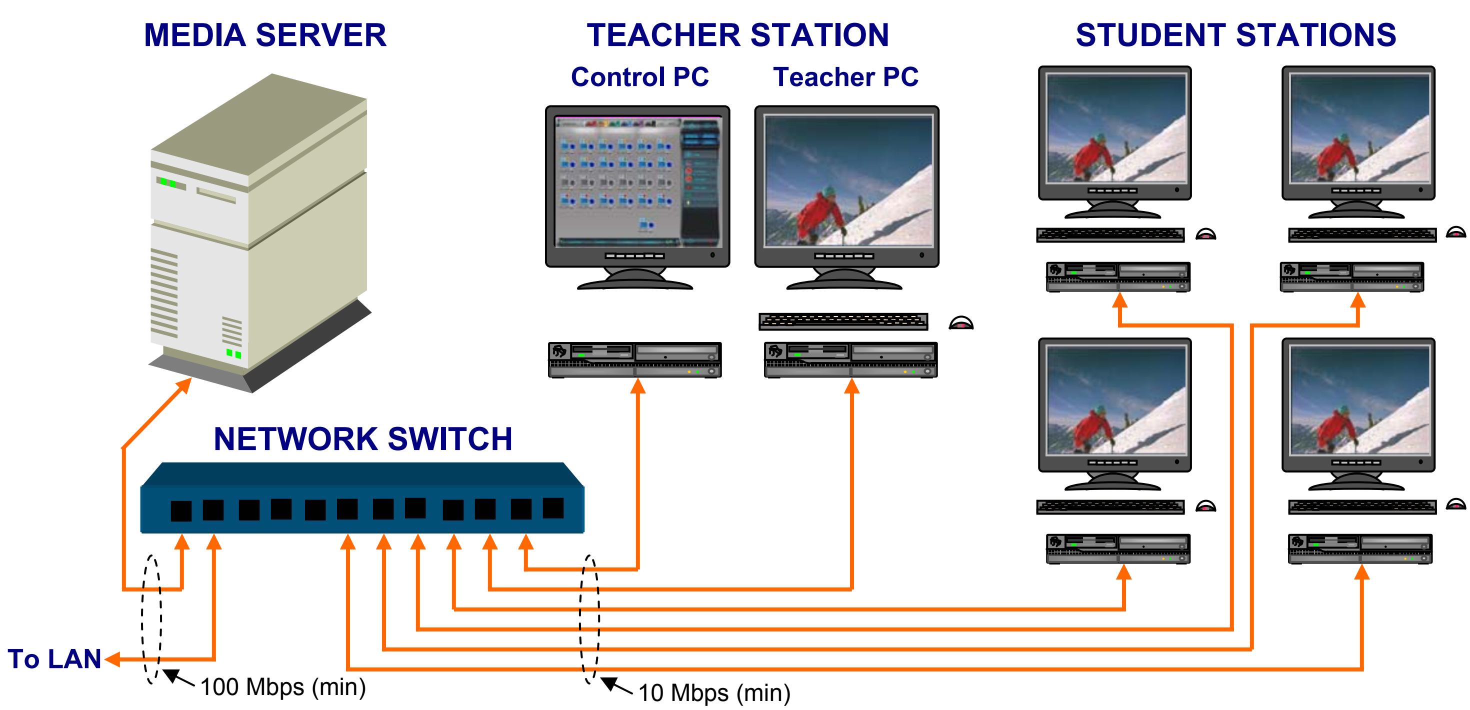

Network Switch

It is strongly recommended that all computers in the language lab be connected to the same network switch to ensure that all communications within the lab are not impacted by external traffic on the LAN. In situations where there are a large number of computers, a *stackable* switch (comprising multiple physical modules) can be used. The minimum bandwidth for the client computers and the control computer is 10 Mbps, although 100 Mbps connections are recommended.

The Media Server should also be connected to a port on the same switch. The minimum bandwidth for the server connection is 100 Mbps, although a 1000 Mbps link is recommended for environments having more than 40 client stations.

Typically, the network switch will be connected to the rest of the Local Area Network, (also via a high-speed port), to provide users within the lab with access to the Internet and other network resources such as printers and application servers (e.g. email).

Shared Media Server

In environments where the Media Server is shared by more than one language lab, it may not be practical to connect all stations in all labs to the same network switch. In such a case, great care must be taken to ensure that each lab has high-bandwidth access to the server, and that external applications do not interfere with client-server network traffic.

Network Protocol

Symposium's client-server communications require a network protocol called TCP/IP, (Transmission Control Protocol/Internet Protocol). The LAN must support this protocol.

#####