PACS 101

Things You Need To Know When Embarking On A PACS Project

Executive Overview

PACS – four little letters that can instill fear in the savviest of radiology administrators. With these four little letters have come the promises of improved patient care and lower costs. Along with these promises have come numerous failures in PACS projects, prompting many to question the viability of embarking upon a PACS implementation.

Many of the issues that you are going to be going through in your PACS implementation are similar to those you are currently facing in the film world. There are various drivers in the film world that require things to be done a certain way. Take for example storage of patient film jackets. Many facilities will store a certain number of months or years of film jackets in their on-site film library, along with other older associated jackets as needed for ongoing assessment of the “current” patients. Older, inactive jackets are usually sent to an offsite storage facility where they can typically be accessed within a few days as needed. Why do we have this dual storage method? It is quite often a matter of available space or the cost of space. Space in a medical facility is usually very expensive, so you do not want to utilize any more space than what is absolutely necessary to store film jackets over the amount that is needed to effectively provide your imaging and diagnostic services. The same is true with digital storage of medical images. It would be great if we could write everything out to the highest speed hard drives and always have instantaneous access to all of the patient images, but this would be cost-ineffective, just as storing all of the films at your main imaging facility is not cost effective.

As you embark upon your project, always be sure to remember that this new technology is not magic. Some vendors will lead you to believe that they have the magic beans that will make all of the implementation issues magically disappear. There are no magic beans; no magic hardware that will make electrons move faster through one vendor’s network line than another’s; no magic compression algorithm that will drastically reduce the image data set size without loosing the original image detail. Everyone is using the same general types of hardware and compression algorithms, with everyone being required to obey the same laws of physics. The differences between the vendors will be the specific methods they use to bring all of these items into a useable solution.

It seems as if one of the main issues that confront business managers is the technology of PACS, and how it all fits into an ROI analysis for the system. What are the issues that need to be taken into account during your analysis? Without at least a cursory understanding of the technology, and the new issues that technology may bring to the table, how can you be expected to account for them during the implementation of the system? Do I have the necessary personnel and infrastructure to allow effective implementation of a film-less system? Regardless of how large or small of an implementation is in the plan, you must understand the issues you are trying to address, and also understand how the technology can assist in overcoming those issues. Where is the technology going to potentially cause issues? Installing new technology without a clear understanding and vision of the business requirements is a perfect equation for failure.

When designing your infrastructure, it is also important to define the requirements around fault tolerance and business continuity. The question boils down to “What happens if something fails?” It is the same types of considerations that you looked at when implementing your film-based environment. If you put in a single film processor, and that processor breaks, how do you continue to image your patients? Typically, you either waited until the maintenance person fixed the unit or you had a 2nd processor that could be used. If you had to wait for maintenance, you were at the mercy of the your service people. Patients might have to be sent to other facilities to have their imaging done until things were working again, causing you to loose revenue. If you had a second processor, there are the associated costs of buying and maintaining 2 units. You could buy 2 smaller, less expensive processing units whose combined throughput could provide you with your total throughput requirements, but in the event of a failure of one unit your throughput capabilities would be effectively cut in half and increase patient wait time. If you bought 2 processors, each capable of supporting the full throughput, the initial cost and ongoing maintenance would typically be higher, but you would be covered in the event of failure of a single unit. These issues are no different with PACS hardware. The same analysis must be done for each item in your planned PACS infrastructure. Make a detailed drawing of everything in your system: data acquisition devices (CR, CT, MRI, film digitizers, etc.), servers, storage devices, networking infrastructure (routers, hubs, switches, bridges, WAN lines, etc.). Once you have this diagram, start crossing things out and ask yourself “What impact would failure of this unit have on my ability to continue to effectively do business?” Remember, if your plan is to go to a film-less environment, you will not have the physical film to fall back on. If you are planning to use film as your backup, be sure that there is not infrastructure within your facility whose failure would also render you not able to produce films. For example, in the past it was common for modalities to be connected by a direct connection to the laser printer for filming. Now, many modality manufacturers are using DICOM printing, where the data is transmitted over the network to the printer. If a network switch were to fail, that switch could be supporting both the PACS system and the film printer, rendering both inoperative and leaving you without your backup infrastructure.

Another thing that you need to assess and plan for is the “ease of access” factor. It is not uncommon for a project manager to base their calculations of image data access based upon the manual methods that are currently used in the film environment, such as the number of times that a film jacket is pulled. When you implement a PACS system, access to the image data is greatly improved, thus making it easier for care providers who may have a need to see those images to now open them up and use them in development of a patient care plan. It is important that you try and take this factor into account, as it will place a greater demand on the infrastructure used to drive the imaging system. If you size and implement a system for use by the radiologists only, then try and expand the access to hundreds of clinicians and referring physicians, you could find yourself needing major system upgrades to maintain the performance demanded by your core business.

The purpose of this document is to delve into the technology of PACS, and to hopefully provide you with a tool to better understand the what’s, why’s and how’s of the technology, along with other issues that you should keep in mind as you begin your assessment process for the imaging project. We will start with the basics to insure a strong foundation in computer and imaging fundamentals. We will then apply these basics to PACS and medical imaging.

This page is intentionally blank ……

Table of Contents

1 Imaging Basics

In order to truly understand all of the issues associated with an imaging project, we must start with the basic building blocks for both computers and digital images. These basic building blocks are: 1) Bits & Bytes and 2) Pixels. By effectively understanding these two items, the remainder of the technology discussions for image acquisition, storage and communication should become much more clear.

1.1 Bits & Bytes

1.1.1 Bits

Computers operate using the base-2 number system, also known as the binary number system (just like the base-10 number system is known as the decimal number system). Computers use binary numbers, and therefore use binary digits in place of decimal digits. The word bit is a shortening of the words "Binary digIT." Whereas decimal digits have 10 possible values ranging from 0 to 9, bits have only two possible values: 0 and 1. Therefore, a 4-bit binary number is composed of only 0s and 1s, like this: 1011.

With binary numbers, each bit holds the value of increasing powers of 2. That makes counting in binary pretty easy. Starting at zero and going through 15, counting in decimal and binary looks like this:

Decimal		Bit 4 (x * 2^3)	Bit 3 (x * 2^2)	Bit 2 (x * 2^1)	Bit 1 (x * 2^0)		4-Bit Binary
0	=	0	0	0	0	=	0000
1	=	0	0	0	1	=	0001
2	=	0	0	1	0	=	0010
3	=	0	0	1	1	=	0011
4	=	0	1	0	0	=	0100
5	=	0	1	0	1	=	0101
6	=	0	1	1	0	=	0110
7	=	0	1	1	1	=	0111
8	=	1	0	0	0	=	1000
9	=	1	0	0	1	=	1001
10	=	1	0	1	0	=	1010
11	=	1	0	1	1	=	1011
12	=	1	1	0	0	=	1100
13	=	1	1	0	1	=	1101
14	=	1	1	1	0	=	1110
15	=	1	1	1	1	=	1111

When we count with binary numbers, the total number of bits used will determine the maximum value that can be represented. The following table represents the minimum and maximum values that can be represented by binary numbers ranging from one bit up to eight bits:

Binary Size		Binary Min	Binary Max	Decimal Min	Decimal Max
1 Bit	=	0	1	0	1
2 Bit	=	00	11	0	3
3 Bit	=	000	111	0	7
4 Bit	=	0000	1111	0	15
5 Bit	=	00000	11111	0	31
6 Bit	=	000000	111111	0	63
7 bit	=	0000000	1111111	0	127
8 Bit	=	00000000	11111111	0	255

1.1.2 Bytes

Single bits are rarely seen alone in computers. They are almost always bundled together into 8-bit collections, and these collections are called bytes. Bytes are, for example, used to hold individual characters (8-bits per character) in a text document or for pixel data in an image.

Why are there 8 bits in a byte? A similar question is, "Why are there 12 eggs in a dozen?" During the development of computer technology, developers decided to group the computer bits into 8-bit collections to improve how the internal workings of the computer could move data around more efficiently.

In medical imaging devices (CT, MR, US, CR, etc.), you will typically see references to 8-bit, 10-bit or 12-bit grayscale images. What this means is that each digital element of the image (called a pixel, to be discussed in a later section) will be represented by 8-bit, 10-bit or 12-bit values. That gives each pixel the following value ranges:

Binary Size		Binary Min	Binary Max	Decimal Min	Decimal Max
8 Bit	=	00000000	11111111	0	255
10 Bit	=	00000000 00000000	00000011 11111111	0	1,023
12 Bit	=	00000000 00000000	00001111 11111111	0	4,095

Note that for the 10-bit and 12-bit images we had to use 2 bytes of data to represent the values. The computer does not break things down below the size of a byte, therefore the upper 6 bits of the 10-bit value and upper 4 bits of a 12-bit value are zero.

To complicate things even further, when we discuss color medical images (US, Nuc Med, etc.) we will talk about 24-bit (and even 32-bit) images. In these images, each color plane (red/green/blue) will be represented by 8 bits of data, effectively tripling or quadrupling the amount of space required to store the image.

One caveat to the bit verses byte discussion is when we look at computer network communications. When computers are connected to network infrastructures, data is sent out on the line by bit via a network interface card. We will discuss this in more detail in a later section.

1.2 Pixels

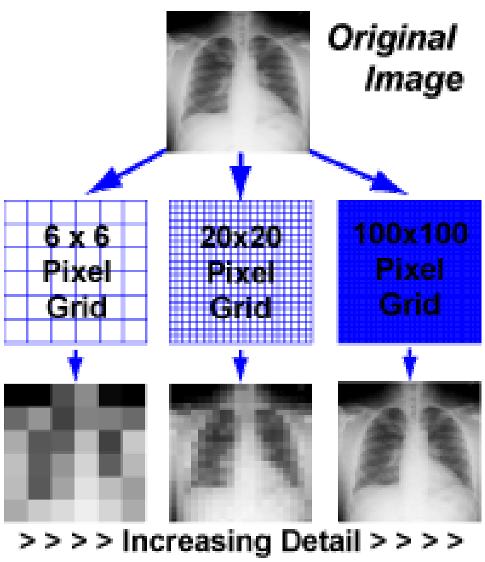

The word pixel stands for “picture element”, and represents the smallest element in a digital image. A digital image is a 2-dimensional array of pixels that represent an image. As we discussed in the previous section, the intensity value of each pixel will be represented by a certain number of bits (stored in complete bytes) of data. Think of your array as a piece of graph paper that is placed on top of your image. Each individual grid element is a pixel, and will have a value based upon the density of the original image at that location.

For example, if we talk about a 640x480 pixel image, it means that the image consists of 640 rows of pixels, with 480 pixels in each row. This would equate to a total of 307,200 total pixels (640x480 = 307,200), each having a separate grayscale or color intensity value. This array size is also known as the spatial resolution. As we increase the size of the pixel matrix, we increase the total number of pixels in the image.

For example:

Note : Shaded cells represent the most common resolution/bit depth configurations.			Storage Size (Bytes)			
Array Size		**Total Pixels**	**8-Bit Image (1 Byte/Pixel)**	**10,12&16Bit Images (2 Bytes/Pixel)**	**24-Bit Image (3 Bytes/Pixel)**	**32-Bit Image (4 Bytes/Pixel)**
256x256	=	65,536	65,536	131,072	196,608	262,144
512x512	=	262,144	262,144	524,288	786,432	1,048,576
1,024x1,024	=	1,048,576	1,048,576	2,097,152	3,145,728	4,191,304
2,048x2,048	=	4,194,304	4,194,304	8,388,608	12,582,912	16,777,216
2,048x2,560	=	5,242,880	5,242,880	10,485,760	15,728,640	20,971,520
4,096x5,120	=	20,971,520	20,971,520	41,943,040	62,914,560	83,886,080

Notice that when you double the size of the matrix (e.g. 512x512 to 1,024x1,024), you have effectively increased the number of pixels by a factor of 4.

As you can see, the numbers get rather large. Because of this, you will typically see the image sizes represented with various prefixes like kilo-, mega- and giga-, as in kilobyte, megabyte and gigabyte. These prefixes are typically shortened to KB, MB and GB. The following table shows the multipliers for each prefix:

Prefix		Multiplier
Kilobyte (KB)	=	1,024
Megabyte (MB)	=	1,048,576
Gigabyte (GB)	=	1,073,741,824
Terabyte (TB)	=	1,099,511,627,776

You can see in this chart that kilo is about a thousand, mega is about a million and giga is about a billion. So when someone says, "This computer has a 2 gig hard drive," what he or she means is that the hard drive stores 2 gigabytes, or approximately 2 billion bytes of data. How could you possibly need 2 gigabytes of space? When you consider that one average 14”x17” CR image is approximately 10 megabytes, you can see that it would only take 200 full resolution images to fill the 2GB of space. In many high-volume facilities, that does not even represent a full day of exams. Although it may sound massive, multi-Terabyte databases are common these days for high-volume medical imaging facilities.

1.3 Mixing of Units

One thing to be particularly aware of is the difference in the units when you are looking at the various components of your imaging system. RAM, hard drives, disks, tapes and other data storage devices will typically represent their capacities in ***bytes*** (Kilobytes=KB, Megabytes=MB, Gigabytes=GB and Terabytes=TB). For networking and data communication lines, the data rates are expressed in ***bits*** per second (Kilobits/sec=Kbps, Megabits/sec=Mbps, Gigabits/sec=Gbps). When you are doing your calculations for transferring data across a network or communication line, be sure to convert your bytes of data into bits, or you will be sizing your lines 8X to small (again, 8 bits/byte).

1.4 Summary of the Basics

You should now have a better understanding of how data is represented on a computer and the basics of how pixels are used to build digital images. In the remainder of this document, we will be relating them items back to PACS technology, and what it all means in regards to your imaging project.

2 Image Acquisition

The first step in the process of medical imaging is acquisition of image data. In this process the image will be acquired in a way that it can be stored and viewed by the system users. We generally consider 3 main methods for acquisition of the medical images:

- Digitization of standard radiographs
- Digitization of the video from the console of the modality, called video frame grabbing
- Image data is captured digitally at the modality (CT, MR, CR, DR, US, DF, etc.)

2.1 Digitizing Radiographs

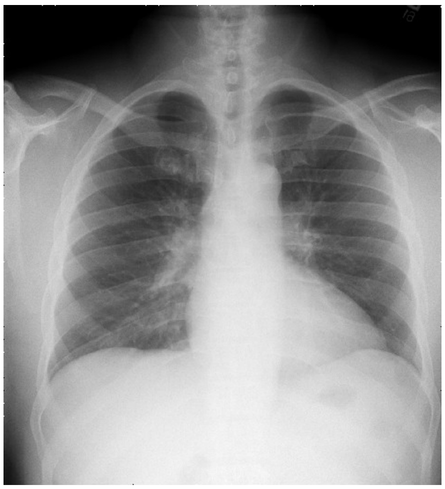

In many facilities, film is the medium for recording all of the medical images from all sources for primary diagnostic interpretation. Film digitization can provide a common, cost effective method to use for acquiring the images for use in a PACS or teleradiology application, although there could be limitations that may hamper its use in a primary diagnostic role.

Whether the films are radiographs produced using standard film/screen techniques or laser printed films of modality images (CT, MRI, US, etc.), a film digitizer can be used to convert these film images into digital images. It is important to understand that the image created by the film digitizer will be a representation of the image or images that are on your original film. If you have scanned a film with multiple CT or MR images, you will not have the ability to make modifications to a single image that was on that film.

Once in the digital format, the film image(s) can be sent to remote locations via network or telecommunication lines.

The film digitization process is a very straightforward procedure. It is a process by which the film image will be sampled into a two dimensional array of pixels. During the digitization process each pixel will be given a grayscale value based upon the density of the image at that location on the film. The size of the resulting uncompressed digital image file will be primarily determined by 2 factors: 1) the size of the pixel data array; and 2) the number of shades of gray that will be represented for each pixel.

2.1.1 Sampling Rate

When digitizing film, the film digitizer unit controls the sampling rate (spatial resolution) at which the film will be scanned, and will create the pixel data array based upon this sampling rate. The most typical sizes of the pixel data array for digitized films will be either 1,000x1,200 pixels (normally referred to as a 1K scan) or 2,000x2,500 pixels (normally referred to as a 2K scan). To use the graph paper example again, think of the sampling rate as the size of the grid on the graph paper, with the larger scan rate equating to a smaller grid element size on the graph paper. As we increase the sampling rate, we are increasing the level of fine detail that will be captured from the film, but as we previously discussed, we also create a much larger image data array.

Sampling Rate Example: The following images show the effect of sampling an image at different rates. The original image on the left was created at a resolution of 100x100 pixels. Each of the re-sampled images was created from the original image using the array size specified. Notice that as we decrease the sampling rate the effective pixel array size is smaller, meaning that it could be transferred faster, but we loose detail in the image that could be critical to making a primary diagnostic interpretation.

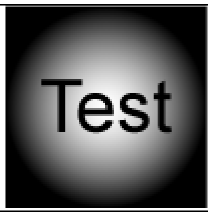	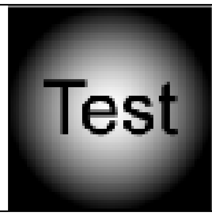	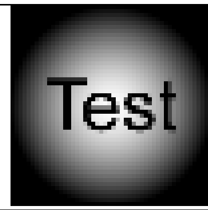	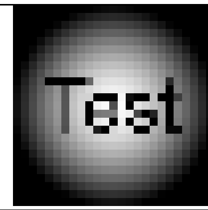	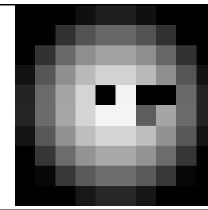
Original Image 100x100 Pixels 10,000 Total Pixels	**Re-sampled Image 75x75 Pixels 5,625 Total Pixels**	**Re-sampled Image 50x50 Pixels 2,500 Total Pixels**	**Re-sampled Image 25x25 Pixels 625 Total Pixels**	**Re-sampled Image 10x10 Pixels 100 Total Pixels**

2.1.2 Grayscale Resolution

During the film scanning process, the digitizer also controls the number of unique shades of gray that can be represented by each pixel. Most digitizers typically support both 8-bit and 12-bit grayscale resolutions. As we decrease grayscale resolution from 12-bits (2 Bytes) to 8-bits (1 Byte), we are decreasing the level of subtle details that we will capture from the film, but the file size that has to be transferred is half the size due to the storage of 1 byte of data per pixel instead of the 2 bytes required for the 12-bit image.

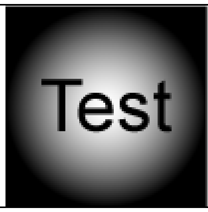	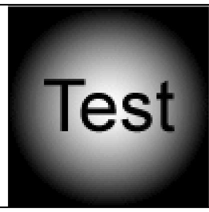	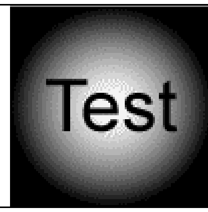	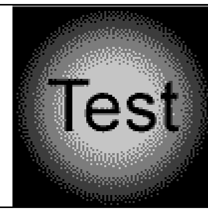	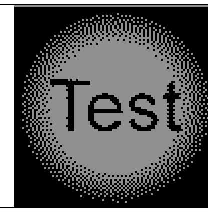
Original Image 100x100 Pixels 8 Bit Grayscale	**Original Image 100x100 Pixels 6 Bit Grayscale**	**Original Image 100x100 Pixels 4 Bit Grayscale**	**Original Image 100x100 Pixels 2 Bit Grayscale**	**Original Image 100x100 Pixels 1 Bit Grayscale**

2.1.3 Sample Spatial Resolutions and Grayscale Resolutions

	100x100 Pixels 10,000 Total Pixels	**75x75 Pixels 5,625 Total Pixels**	**50x50 Pixels 2,500 Total Pixels**
8 Bit	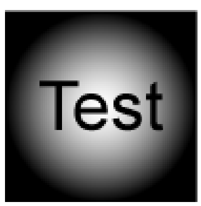	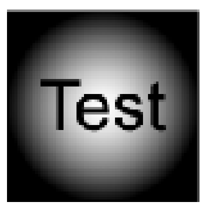	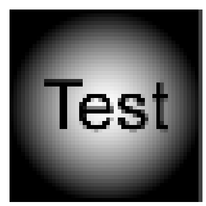
4 Bit	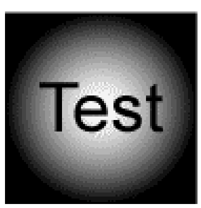	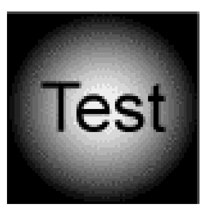	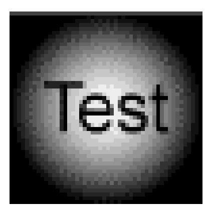
2 Bit	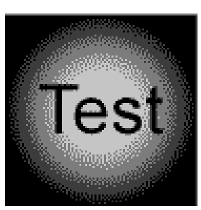	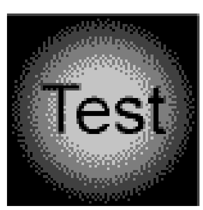	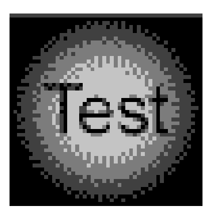

2.2 Types of Film Digitizers

There are generally 3 types of film scanning technologies available on the market to digitize films, with each having positive and negative factors to be considered: 1) Camera-On-A-Stick; 2) CCD Film Scanners; and 3) Laser Film Scanners.

2.2.1 Camera On A Stick

This technology uses a digital camera on a mounting device (some as simple as a standard tripod) and pointed at a film positioned on a light box. The camera is used to take a picture of the film using its internal setup for the pixel data array size and the bits per pixels. This method is typically the least expensive, but also yields the poorest overall image quality. The most typical quality issues arise from poor setup of focus and aperture settings at the camera, poor illumination of the film and dirt on the lens. The quality issues associated with this method have made this an uncommon choice for most medical users outside the veterinary world.

C-O-A-S Pros		C-O-A-S Cons
▪ Very Low Cost		▪ POOR Image Quality

2.2.2 CCD Film Scanners

CCD film scanners typically use either a fluorescent bulb or red LED as a light source. The light that passes through the film is read by a row of solid-state CCD detectors, with the output converted into a numeric value representing the pixel's grayscale. The film is fed through the scanner automatically, with full rows of data captured simultaneously as the film moves through the scanning unit. These systems are more expensive than the camera-on-a-stick systems, but less expensive than a laser scanner unit. Image quality also falls into the middle of the spectrum. The biggest image quality issues are pixel "bleeding", caused by light scattering due to pixel-to-pixel overlap, and a lower effective grayscale dynamic range (dark regions do not turn out as dark in the digital image file) due to the light strength at the illumination source. System maintenance on these units is typically lower. Some units can be configured with an automatic sheet feeder to facilitate ease of digitizing multiple films for a single patient.

CCD Pros		CCD Cons
▪ Moderate Cost		▪ Pixel "Bleeding"
▪ No detector drift		▪ Grayscale Dynamic Range

2.2.3 Laser Film Scanners

Laser film scanners use a high intensity laser as an illumination source. The laser light that passes through the film is read by a photomultiplier tube and converted into a grayscale value. The film is fed through the scanner automatically, with full rows of pixel data captured as the laser is moved across the film using a galvanometer. These systems are typically more expensive than the CCD systems, but also provide the highest image quality of all the available digitizing technologies. System maintenance on these units is typically higher, as there are more moving parts that can fail or go out of calibration. Some units can be configured with an automatic sheet feeder to facilitate ease of digitizing multiple films for a single patient.

Laser Pros		Laser Cons
▪ Best Image Quality		▪ Higher Cost
		▪ Higher Maintenance

2.2.4 American College of Radiology Digitizer Recommendations

For primary diagnostic needs, the American College of Radiology (ACR) recommends:

- Spatial resolution of 2.5 line pairs/mm (lp/mm), which equates to approximately 1,800x2,200 pixels for a 14"x17" film or 1,000x1,250 pixels for an 8"x10" film.
- Grayscale resolution of 10 bits per pixel, providing 1,024 shades of gray.

There are no film scanner system recommendations for non-primary diagnostic uses.

2.3 Video Frame Capture

The next method for image acquisition is to digitize the video as it is displayed on the console of a modality. This process is typically referred to as a frame grab, and is done in the same way as images are captured to send to film output. This method is commonly used to support older modalities that do not support DICOM.

When a modality is doing a study, the technologist will press a button on a keypad signaling the Video Capture host PC to capture the video signal that is currently coming from the display on the modality. The video data will be captured into an image data array like that discussed in the film digitizer section. The spatial resolution of the captured data is dependent upon the modality, but are generally as follows:

Modality		Typical Frame Grab Image Resolutions
US		512x512x8bits (If color images, 8bits/color, 3 Bytes)
CT		1024x1024x8bits or 512x512x8bits
MRI		512x512x8bits
NucMed		512x512x8bits (If color images, 8bits/color, 3 Bytes)

When capturing data via video frame capture, it is important to note that the images will contain no sizing or positional data, therefore you cannot measure in real units, use scout lines or do overlays. Also, to see multiple window and level settings for an image, you must grab the images for each W/L setting at the modality, and transmit each individual image. Adjustments at the workstations will not be working with the true window and level values, so you would really be adjusting the brightness and contrast of the frame grabbed images.

2.4 DICOM Image Acquisition

The last method for acquiring images is via a direct DICOM transfer of the original image data from a device (CT/MR/US/DF/CR/DR/WS/etc.). This method will typically provide the highest image quality and greatest flexibility, as you are using the original data as is created at the device. This will allow you to perform functions such as changing window and level settings on the original data set and doing measurements in real units for images sent with pixel sizing data from the modality. This type of data can be critical in a primary diagnostic interpretation scenario.

Modality		Typical DICOM Image Resolutions
US		512x512x8bits (If color images, 8bits/color, 3 Bytes)
CT		512x512x12bits
MRI		512x512x12bits
CR - Fuji		2,000x2,000x10bits 35x35 image size
CR - Kodak		2,000x2,500x12bits 35x43 image size

2.5 Image Acquisition Summary

When looking at the specific methods that will be used for image acquisition, you must closely consider how the images are going to be used. What is the intended use of the system? Will it be a basic teleradiology system used for emergency night over-reads only, hence you are looking for gross abnormalities in the images? Are you trying to move your facility to a film-less environment where it will be critical to provide the radiologists with the highest image quality possible for primary interpretation? Are you looking at a cost effective way to transfer images for primary diagnosis from a remote hospital or clinic? Let the intended use of the system and its inherent image quality requirements drive the image acquisition methodology. Don't let the tail wag the dog.

3 Data Communication

Once you have acquired the image data, you now have to transfer that data to one or many destinations within your facility. There are going to be a number of communication issues that you have to take into account when looking at this data communication process.

3.1 Local Area Networks (LAN)

When looking at implementing your PACS or teleradiology system, be sure to take a close look at your plans for networking the imaging equipment, especially if you plan to use existing networking infrastructure. As you can tell by the previous discussions in this document, medical imaging applications have the potential to create a huge amount of data. Depending on your implementation, all of this data will be communicated across your network an average of 1-5 times before the final diagnosis. These transfers would include the initial data transfer from the image source to the image server and one to many pulls of the images to workstations for the interpretation. Be sure not to leave out data transfers related to any automatic pulling of historical comparison studies because depending on the type of study, this could easily double the amount of data to be transferred.

Another issue to take into account is peak traffic periods. In many medical imaging operations, studies are not spread out linearly throughout the day, therefore the demands placed on the network to transfer the data from those exams is not linear. There are going to be peak loads that must be understood and planned for. Do not make the mistake of calculating your network bandwidth requirements based upon the average number of anticipated studies for the day. Your PACS system is suppose to allow you to be more efficient and do more studies. If the network cannot take the peak loads (both current peak loads and anticipated peak loads based upon expected growth), you will cause bottlenecks in the system and erase any potential for increased efficiency.

It is critical that you do an in-depth analysis of ALL data transfer requirements to insure that you provide adequate network bandwidth for the expected volumes of data. Without this analysis, you are almost guaranteed to run into performance problems with your imaging application and any other applications that you try to run across the same network infrastructure. For high-volume applications, it will be critical that you segregate the imaging data network traffic onto its own high bandwidth infrastructure.

3.2 LAN Speeds

When looking at current network topologies, we will find that there are generally 3 topologies in use:

- 10 Mbps Ethernet
- 100 Mbps Ethernet
- 1,000 Mbps (also know as Gigabit Ethernet)

Each of these network topologies provides a finite amount of bandwidth available for use in transferring data. It is important to understand that regardless of the topology you use, you will not get 100% utilization of your network bandwidth. It is best to use a maximum utilization rate of approximately 80%. Above this level, data collisions will begin to degrade overall network performance at a dramatic rate.

The specific hardware used in the implementation of the network (switches, hubs, routers and bridges) can also have a dramatic impact on your performance. Some network hardware devices, such as smart switches, provide data segregation between each node of the device, effectively increasing the overall bandwidth available to the computer devices connected to the network. Other hardware can support differing packet sizes that, especially on the Gigabit Ethernet networks, can dramatically improve the system performance.

If you are planning to utilize existing networks, but sure to get a clear handle on how each device operates. If you are purchasing new network components on which your PACS system will operate, be sure to invest in the devices that will give you the needed performance. Realize that the small amount you might save by going with a cheaper, less-capable device will haunt you down the road, and typically sooner than later.

When you start to look at a large network spread out across a big facility, you will inevitably have numerous network devices that the data must pass through for a transfer. Each network device adds latency to the overall transfer and can dramatically impact the end-to-end performance. Also, it is not uncommon to have multiple network topologies in a large network, with high-speed topologies making up the “backbone” of the network communications and the slower topologies running to the desktop. In these layouts, be sure to understand where these potential bottlenecks are and be sure not to layout your PACS devices such that the image data must routinely traverse this portion of the network.

3.3 Wide Area Networking (WAN)

In most medical imaging system implementations, there will be requirements to transfer data across WAN communication lines. Whether you are using T-1s, DSL or any other type of communication lines, this represents a probable bottleneck for data transfers. Be sure to factor in those peak loads that we discussed earlier, or you will be very disappointed with your system’s ability to handle the increased loads during your busy times of the day.

You can typically address communication line bandwidth issues in one of two ways: 1) Add additional lines to get additional bandwidth or 2) Compress the data to reduce the effective amount of data that must be transferred across the line. As with most options, each solution comes with trade offs.

If you make the decision to add communication lines, this will cost you more money for the increased bandwidth. If you have sized your communication lines based upon peak loads (as you should) and the peak loads are a drastic swing from the steady state, this bandwidth may sit relatively idle for the rest of the day.

If you decide to use compression algorithms to reduce the amount of data that must be transferred, the biggest trade off will typically be in image quality. It really depends on how undersized the communication lines are in relation to the amount of uncompressed data that you have to transfer. Compression ratios of approximately 3:1 are considered to be lossless which means that you decrease the amount of data by a factor of 3, and can recreate the original image data during the decompression process after the transfer. If there is a large disparity in the amount of available bandwidth in relation to the amount of uncompressed data to be sent, you will be required to use a large compression ratio. This will have a direct impact on image quality, as the more you compress the more subtle detail within the image is lost. If you are making a primary diagnosis from the images, a large compression ratio could impact image quality to a point that accuracy in the interpretation could be affected.

Also, with the use of WAN communication lines you may be faced with the requirements to use secure protocols to insure patient data confidentiality as is demanded by HIPAA regulations. These secure protocols use different types of packets transferred across the WAN, and will incur additional transfer overhead of approximately 5-10%. Some Internet Service Providers (ISPs) have hardware that does not support these secure protocols. If you are using an ISP for your Internet access, and will be using these communication lines in your imaging system, be sure to verify that they can support the use of the secure protocols (VPNs or IPSec).

A final word of caution; be very careful in the use dialup phone lines. These communication lines will only provide a maximum of ~50Kbps throughput, forcing you to use very high compression ratios. This coupled with the overhead of creating a secure connection (VPN) will make the performance and image quality issues large obstacles to overcome.

3.4 Networking Examples

The best way to get your hands around these issues is to work the math on some examples. In this way, you can begin to gain an appreciation for what does and does not work. In all of these examples, we are assuming that the local or wide area network line is 50% utilized by other applications, except for dialup lines, which typically cannot be shared simultaneously between multiple users.

3.4.1 20MB Data Set: (2) FD/CR Images - 2Kx2.5Kx2Bytes

Local Area Network - 50% Utilized

Network Speed (Mbps)	Avail. Bandwidth (Mbps)	Latency (Sec) Network	Latency (Sec) Comp/ Decomp	Transfer Speed (Sec) - 20MB Dataset @1:1 Ratio	@3:1 Ratio	@8:1 Ratio	@20:1 Ratio
10	5	2	2	42.00	17.33	9.00	6.00
100	50	2	2	6.00	5.33	4.50	4.20
1,000 Standard Packets	217	2	2	2.92	4.31	4.12	4.05
1,000 Jumbo Packets	500	2	2	2.40	4.13	4.05	4.02

Wide Area Network - 50% Utilized

WAN Speed (Kbps)	Avail. Bandwidth (Kbps)	Latency (Sec) Network	Latency (Sec) Comp/ Decomp	Transfer Speed (Sec) - 20MB Dataset @1:1 Ratio	@3:1 Ratio	@8:1 Ratio	@20:1 Ratio
1,544 T-1	772	2	2	261.07	90.36	36.38	16.95
512 DSL/Cable	256	2	2	783.25	264.42	101.66	43.06
256 DSL/Cable	128	2	2	1564.50	524.83	199.31	82.13
53 Dialup	53	2	2	3775.58	1261.86	475.70	192.68

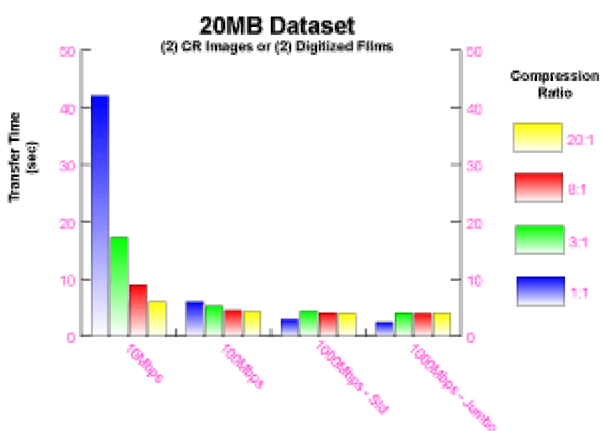

LAN

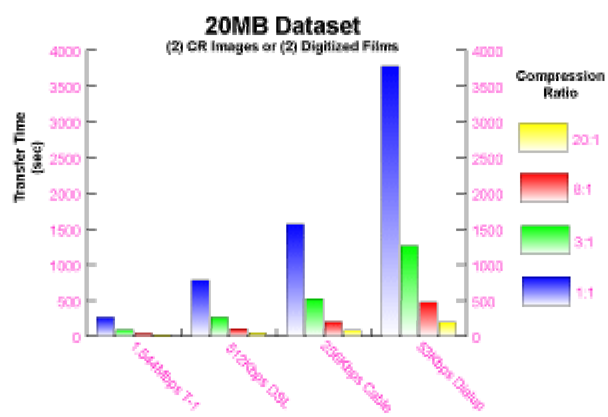

WAN

3.4.2 63MB Data Set: (120) CT Images - 512x512x2Bytes

Local Area Network - 50% Utilized

Network Speed (Mbps)	Avail. Bandwidth (Mbps)	Latency (Sec)		Transfer Speed (Sec) - 63MB Dataset			
		Network	Comp/ Decomp	@1:1 Ratio	@3:1 Ratio	@8:1 Ratio	@20:1 Ratio
10	5	2	2	128.00	46.00	19.75	10.30
100	50	2	2	14.60	8.20	5.58	4.63
1,000 Standard Packets	217	2	2	4.90	4.97	4.36	4.14
1,000 Jumbo Packets	500	2	2	3.26	4.42	4.16	4.06

Wide Area Network - 50% Utilized

WAN Speed (Kbps)	Avail. Bandwidth (Kbps)	Latency (Sec)		Transfer Speed (Sec) - 63MB Dataset			
		Network	Comp/ Decomp	@1:1 Ratio	@3:1 Ratio	@8:1 Ratio	@20:1 Ratio
1,544 T-1	772	2	2	818.06	276.02	106.01	44.80
512 DSL/Cable	256	2	2	2462.94	824.31	311.62	127.05
256 DSL/Cable	128	2	2	4923.88	1644.63	619.23	250.09
53 Dialup	53	2	2	11888.79	3966.26	1489.85	598.34

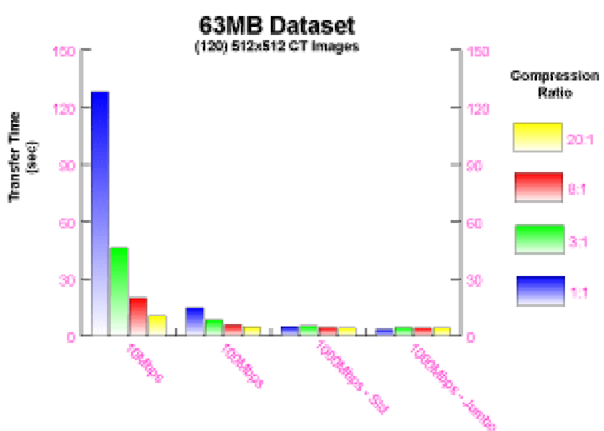

LAN

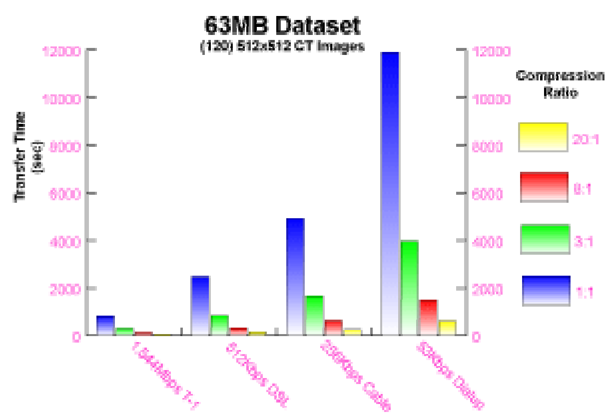

WAN

3.5 Data Communication Summary

Effective LAN and WAN networking can make or break your imaging project. When starting on your implementation, be sure to start with a detailed network diagram that shows all of the network devices within your system. From this diagram, try to identify any initial problem bottleneck areas to address.

If you are going to add new network infrastructure to support your imaging system, be sure to invest in hardware that can run at least the highest common bandwidth available at the time (Gigabit Ethernet at the present time). Even if you initially decide to run the network at a slower speed, you will have the ability in the future to upgrade to the faster speed if/when required.

If you were going to plan to use existing network infrastructure, it would be a wise idea to have your network engineers provide you with an analysis of current utilization of all LAN subnets and WAN communication lines that will be touched by the imaging data. Many times, even the network engineers do not know just how congested a network really is. Without an actual analysis, you will have no guarantee that the bases for your calculations are valid as you embark upon building your imaging system. In this scenario, you may get a very bad surprise when you switch on the imaging system and bump an already busy network into overload and cause issues with other hospital applications sharing the network. Also be aware that your IT department may not have the specialized equipment needed to assess your network utilization. In this case, it is well worth the expense to contract this analysis service to an outside networking vendor. They will have the equipment that can provide analysis over a longer period of time, giving you the ability to see peak load trends that may not be readily evident by looking at the traffic for a single day. If possible, try to get a minimum of one full week of analysis for each subnet and WAN line: two weeks would be even better if possible.

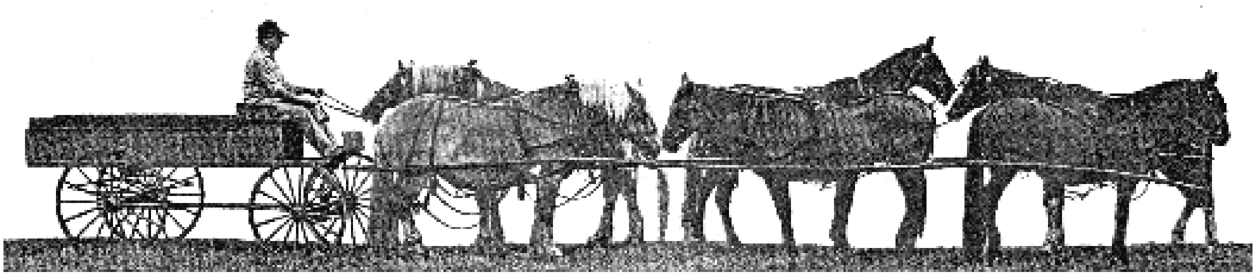

Always remember, there's more to networking than just hooking things up ……

4 Data Storage

Data storage in a medical imaging system can encompass a number of different things:

- Short-term Storage
- Long-Term Storage
- System Backups

The main issue for each of these areas can be quite different, ranging from data access speed to media costs to data longevity to catastrophic event (fire, tornado, hurricane, earthquake, etc.) business recovery. A typical installation would need to have a blended approach to these storage strategies to provide the needed capabilities at the best possible price structure. It is no different than the methodologies you may currently use when storing the physical film jackets, where you have the "short-term" storage area at the main imaging facility, and the "long-term" storage area offsite. It is critical to understand your requirements for these areas, and implement a strategy that will allow you to meet your technical and business objectives.

Again, each of the storage issues will drive tradeoffs that must be considered. When it comes to data storage it is typically a cost verses performance tradeoff, meaning that if you need to have fast access to the data, it is going to cost more for the storage subsystem. This is similar to the cost structure for storing a patient film jacket at your main imaging facility verses an offsite storage facility, where onsite storage will typically provide faster retrieval but will cost more.

Just be sure to understand that there is not a "silver bullet" for storage that can do everything that is needed within a single architecture. In this section, we will discuss each area, and try to give you a better understanding of the potential options and tradeoffs. Additional storage requirements around fault tolerance and disaster recovery will be covered in the Server section, as there will be considerations in this area too.

4.1 Short-Term Storage

Short-term storage is typically defined as the location that all newly acquired images will be saved. It is this area that will be queried by the diagnostic and clinical workstations when images are needed for a patient who had recently been through your facility. Since these are new images there is a guarantee to have to retrieve them at least once and maybe in excess of 5 times in order to provide the initial interpretation. Consequently, this storage area is typically configured with fast-access media (SCSI RAID, SAN, etc.) consisting of very high-speed magnetic hard disks to provide the quickest access to the data as possible.

It is critical that you size this storage area such that it will provide enough space to save patients' images to correspond with the overall length of the treatment plans, including patient follow up. We typically see this range anywhere from 4 weeks to 12 weeks, depending on the types of cases that are seen in the facility.

It is also critical that you take comparison studies into account for your short-term storage requirements. It is great if you have instant access to the new images, but it you need to have access to the old studies in order to make a full diagnosis and have to wait to get that data due to slow long-term media, you are going to be impacting the efficiency of your reading personnel, and will not be able to fully realize the potential of your imaging system.

If possible, workflow methods such as image pre-fetching should be utilized to automatically search the image server for scheduled patient procedures for the next day. When data for a patient is found on long-term storage, this data can be copied from the typically slower long-term media (tape, optical platter, etc.) to the short-term media so it is immediately available for retrieval during the interpretation process for the new procedure to be performed. The pre-fetch process should be kicked off during "slow" times (typically late at night) so there is not an additional load placed on the network during busy operational hours. When doing your calculations for short-term storage space, be sure to look closely at the type of imaging that you do, and how comparison studies are currently used. You can then make a determination as to the "average" number of comparison studies that may be required to be stored. From this analysis, you can size the short-term storage space accordingly.

4.2 Long-Term Storage

Long-term storage can be thought of in a couple of different ways, depending upon the requirements of your facility. The specific methodology you use may be highly dependent upon a number of factors, and will typically employ more than one type of media (magnetic, tape, optical, etc.).

For a small to medium imaging facility, the long-term storage could consist of 2 different medias blended into an effective storage strategy. The price of storage technologies is moving lower over time. This is especially true for magnetic hard drives, resulting in higher-speed, higher-density drives being more cost-effective to deploy. For the small to medium facilities that do not have huge imaging volumes, it could represent a viable alternative to using tape media only for the long-term storage. One could use a cost-effective high-density technology, such as a Network Attached Storage (NAS) device, to archive up to say 5 years worth of data, while also writing the image data out to tape media for offsite storage. This would easily provide sub-minute access to all imaging studies for 5 years, while also providing a tape copy of the data that can be transferred to an offsite storage facility in case of a catastrophic event.

For a large imaging operation, the long-term storage area could be defined as a combination of intermediate storage and long-term storage. The intermediate storage could again use a marginally slower media such as NAS for storage of patient data that is 3-24 months old. This would potentially address the bulk of the comparison study retrievals since studies older than this may not need to be pulled on a routine basis. For the long-term storage of information beyond 24 months, a tape media with an automated media loader/jukebox could be used that would provide a lower cost/study for the storage, but retrieval times would be longer than on the short and intermediate-term storage mediums. The auto-loaded tape media could hold data in the 24-60 month range, with tapes going to shelf storage after that time, requiring manual loading. Again, the tradeoff for the lower cost is longer retrieval times.

4.3 System Backups

In addition to image storage, you will also need to create an infrastructure for backing up the core files (OS, databases, etc.) on your system. Most imaging systems maintain a database that keeps track of patient demographic data and the locations of the image files within the file system (short-term and long-term). These databases typically do not hold the actual image data; rather it stores pointers to where the data resides on the system. On a large system, this database has the potential to grow quite large, potentially exceeding 20GB. It is critical that these database files be backup on a routine basis and taken offsite for secure storage. This will ensure that in the case of a disaster at the facility, the operating system and database files can be quickly restored onto a new system to allow the system to quickly restore operations. If this database backup were not done, the only way to rebuild the information is to load each tape you have created containing images, read the patient demographics from the header of every image file on the tape, and manually recreate the entry for each image in the database. This would have to be done for EVERY tape created by your system, and could easily represent months of work for a large system. If your system were configured with a manual tape drive (no tape autoloader/jukebox), someone would have to manually insert each tape into the drive throughout the restoration process. In contrast, using the backup would allow you to restore your database in a matter of hours from your tape(s).

It is also important to understand the storage requirements for HIPAA compliance. The HIPAA Security Standard, section 164.308, specifies the requirement for medical facilities to maintain data backups that are retrievable exact copies of the data. These standards were effective as of April 21, 2003, with required compliance by no later than April 21, 2005. It is critical to be aware of these compliance requirements, as any system that you implement today will need to provide compliance to these standards in a relatively short timeframe. Working this compliance into your plan from its inception will drastically reduce your headaches as the compliance date approaches.

4.4 Data Storage Summary

Data storage can be a very complex beast to slay. By carefully assessing your requirements during the analysis phase, you can build a cost-effective storage infrastructure that will provide you with the necessary room for planned growth, yet can also be augmented and upgraded to address unanticipated growth within your business. The initial implementation of your storage infrastructure can be difficult to change, especially once you get down the road a ways and have committed a large amount of data to the system. Rebuilding some of the large image storage volumes could take in excess of 16-20 hours, during which time the system would not be available. It is important to get it right from the beginning to insure that you have no impact to system availability.

5 Servers

When we talk about servers, we are talking about the computer system(s) on which the physical image data will be stored and retrieved by the users. These computers are the brains that sit out in front of the data storage devices that we discussed earlier, and direct data movement between the users and the data. Some of these servers may provide basic archival capabilities only, where others may provide expanded features such as web access, HIS/RIS interfaces via HL-7, RIS functionality or other capabilities.

When you are looking at a PACS system, it is critical to understand what is going to be running on the server, the way in which users will be accessing the system and the total number of users that will simultaneously hit the system. These items are typically the drivers that will be used to determine how large of a server you will need based upon the performance your business demands. Again, it is not magic as to how this is determined. As users are accessing the server, they take up resources (CPU time, RAM, network bandwidth, etc.) from the machine. When they log off, those resources are freed up for reuse. By knowing the total number of simultaneous users you expect, you can determine the total server demand that user load will draw, and can configure the hardware components accordingly. Be sure to factor in your planned near-term increases in user loads to insure that you have some headroom.

When designing your server infrastructure, it is also important to define the requirements around fault tolerance and business continuity. There are many different ways in which systems can be designed to take differing amounts of fault tolerance into account. What levels of user intervention is desired in the failover process? How long of a failover time is acceptable? We are back to the cost verses performance discussion on this item. If you require fully automated, real-time fault tolerance, it is going to cost much more than running a system that uses data tapes for backup and could require hours/days to restart operations.

Be sure that you also understand the backup and business recovery requirements in the HIPPA security standards for your facility. The Emergency Operation Plans are included in the HIPAA Security Standard, sections 164.306 and 164.308, and outline the requirements for establishing and implementing business continuation procedures within the covered systems.

5.1 Mirrored Servers

Mirroring servers is the process of setting up and maintaining two identical servers. One unit is typically used as your main server, with all users accessing this system for image data retrievals. The second unit, which is a mirror hardware image for CPU, RAM, storage, etc., is typically placed at an off-site location and connected via WAN communication lines. As the system operates, all image data stored to the primary server is also stored to the second server in what is sometimes called a "cross-committal" of the data. In the event of a failure of server #1, server #2 is manually reconfigured with the IP address of the original primary server, and restarted on the network. This is a manual process, with failover taking 30-60 minutes.

As you can easily tell, the hardware costs associated with the mirrored approach are going to be double that of a non-mirrored system. Also, there will need to be adequate communication line bandwidth available to allow the data to make it to the server in a timely manner so as not to bog down the sending unit. Also, using this cross-committal approach, it could be possible for data to get out of sync, like during times of communication line outage, or even a failure at the secondary server.

5.2 Highly Available Servers

Highly-Available Servers, also known as HA systems, provide the ability to configure hardware that provides internal redundancy and automatic failover capabilities. HA systems will be server computers configured with internally fault tolerant components like hot-swappable power supplies, cooling fans, disk drives, network cards and CPUs that in the event of a failure will automatically failover to a backup. Replacement of the failed component can be done with the system still powered on, and in the case of a hard disk, the new drive would be automatically reformatted and the data rebuilt on the fly.

Each server would be configured with support for all processes required within the imaging system, but would only run a subset of those processes on a day-to-day basis. These server units would be "clustered", so in the event of a failure that takes out one of the servers, the software processes running on that failed machine would be automatically started up by the other server box. This requires the use of additional hardware and software components to allow this fault tolerance to take place in an automated fashion, and will drive the overall cost of the system up. It also important to note that when configuring an HA system, the servers involved are typically in the same room, and therefore are not covered in the event of a local catastrophic event (both units would be destroyed). There are ways to do "wide-area" HA systems, but costs for such systems and the associated communication lines are typically hard to cost justify within a medical imaging system.

5.3 Data Backups, Disaster Recovery and Emergency Operation Plans

Data Backups and disaster recovery plans are all included in the HIPAA Security Standards, which became effective on April 21, 2003. The scheduled date for compliance with these standards is April 21, 2005. Sections 164.306 and 164.308 generally outline the requirements of the standard in relation to patient data security and accessibility.

Data backups as defined by the HIPAA Security Standard are meant to provide a retrievable duplicate copy of the patient data. It does not specifically define a requirement for how long it would take to make this data retrievable, rather it seems that this must be determine by each facility and their specific requirements.

Storage for the purpose of disaster recovery can blur the line with backups and long-term storage. Typically for an effective disaster recovery plan, your data needs to be stored in a geographically dispersed location from the main data source such that the likelihood of a single catastrophic event destroying both copies is unlikely. There is not a clear definition of how far apart "geographically dispersed" means. It is definitely not a location in the same building, or even on the same campus, as there are catastrophic events such as fires and tornadoes that can easily affect a relatively large area. Some might think a few miles away would be OK, but what about hurricanes-prone areas where widespread destruction is not uncommon. There is not a single definition of "geographically dispersed" that will necessarily fit every scenario. It is really up to each facility to make that determination using prudent and reasonable assumptions for their business, and the risk associated with loosing both copies the data they are responsible for. Remember, prior to 09/11/2001 some tenants of the 1 World Trade Center stored their backup data at 2 World Trade Center, thinking that both buildings would never be destroyed.

5.4 Server Summary

As we previously discussed, it is critical to understand the mission-critical nature of the functions provided by each component of your system. It is very important to understand this for the servers, as these are typically the main storage point for the entire network, but you really need to look at this for all components throughout your PACS system. Be sure to create a detailed system diagram. Keep the diagram continuously updated. Keep assessing the impact of loosing the functionality of a device on the overall operation of the imaging system. Can you keep doing business if the unit fails? If not, what will be the financial and business impact of the failure? Is it a manageable risk? If not, which of the fault tolerance solution will fit into your overall business model and cost structure?

6 Summary

We hope that this document has provided you with some additional insight as to some of the things that you will need to take into account during the analysis and implementation of your PACS system. You will be striving to improve upon what you have today in your current environment. All of the issues raised by the newer technology are relatively straight forward IF you know of the issues and take them into account during the analysis phase. You cannot plan for what you do not know; you know that these systems give you the ability to redefine the way in which you can deliver services to your patients, but how and to what extent will these changes be different from the way you do it today?

You now have the ability to put actionable data into the hands of everyone who needs it. The hard part is insuring that the infrastructure that you put in place (hardware, software, etc.) takes into account all of these hands and allows you to improve overall efficiency and patient care.

Good Luck!!