Design Considerations for Medical Image Archive System

Table of Contents

Introduction

Medical images are at the heart of the healthcare diagnostic procedures. They have provided not only a noninvasive mean to view anatomical cross-sections of internal organs, tissues, bone and other features of patients but also a mean for physicians to evaluate the patient's diagnosis and monitor the effects of the treatment, and for investigators to conduct research of the underlying disease.

Medical images come from a board spectrum of imaging sources such as computed axial tomography (CT), magnetic resonance imaging (MRI), digital mammography and positron emission tomography (PET), and they generate a staggering amount of image data and important medical information. Consider a single study of CT scan of thirty 512 x 512 x 16 bits images; each examination is approximately 15 megabytes. And for digital mammography study of four 4K x 6K x 12 bits, the amount of data grow to 200 megabytes per examination. For the Clinical Center, it is estimated that approximately 8 million individual image slices are generated per year. This is equivalent to four terabytes of image data assuming a typical individual image size is 0.5 megabytes.

For a research center, the emphasis may shift from the generation of image to post-processing and data management since a research team may generate even more processed images and other data from the original image after various analyses and post-processing. With all the research information that is generated on a daily basis, the storage requirements and data management are enormous; investigators need a reliable storage system to store, retrieve and preserve their work.

A medical image data repository for health care information system is becoming a critical need. This data repository would contain comprehensive patient records, including information such as demographics, medical history, clinical data and related diagnostic images, and post-processed images. Due to the large volume and complexity of the data as well as the diversified user access requirements, the implementation of the medical image archive system will be a complex and challenging task, in particular with an archive system that is scalable and upgradeable with time.

Currently, the National Institute of Neurological Disorders and Stroke (NINDS) has an immediate need for an image repository to store and to make available a large amount of MRI human stroke image data to investigators and the research community. This collection of MRI stroke clinical trial image data is estimated to require 60 gigabytes of storage space, which include125 participants of approximately 200 to 500 megabytes of data per patient. And over the next two years, the number of participants will grow to 1200 to 1500 and approximately one terabytes of storage space. The National Institute on Aging (NIA) has a near future need for storing a large amount of images and analytical data from the AGES study that will be conducted in Iceland. This study will involve in the collection of data of approximately 10,000 persons.

This paper discusses a proposal to the development of a centralized medical image archival repository in support of medical researchers to analyses medical imaging data from clinical studies. The motive is to realize the greatest possible benefit from gathering, indexing, communicating, managing, and archiving multimedia imaging data that are already exist or being gathered for conducting research, teaching, training and measurement in order to provide a cost effective healthcare delivery, medical education, and clinical research.

The design and consideration of a data storage architecture is focused on the approach of two existing and emerging technologies or models: a) Picture Archiving and Communication System, b) Healthcare Application Service Provider.

PACS

Picture Archiving and Communication System (PACS) consists of various electronic components and connects to a computer network to facilitate the management of medical images from various imaging modalities. Other functions included store and retrieve, manipulate and display of medical images on a network. It also integrates the image management system with radiology and clinical information systems.

PACS consists of image acquisition, archiving and retrieval, communication, image processing, distribution, and display of patient information. Its system infrastructure consists of hardware components such as: imaging device interfaces, storage devices, host computers, communication networks and display system that are integrated by network system and software systems for communication, database management, storage management, and network monitoring. The system offers an efficient means of viewing, analyzing, and documenting study results, and furnishes a method for effectively communicating study results to the referring physicians.

Film has been used to capture, store, and display radiographic images in hospital. Today, many radiological modalities generate digital images that can be viewed directly on monitor displays. Many experts concede that it is only a matter of time before PACS replaces such traditional system. Digital archival will eliminate the need for bulky film room and storage site and the possibility of lost films. PACS has demonstrated an increase in department and hospital productivity by electronically managing digital image data. The system provides an efficient means for archiving, retrieving, and displaying of digital data, and has three major advantages over traditional hardcopy/film-based reading and storage.

1. Enable medical image and data distribution throughout computer networks
2. Enable electronic archiving and retrieval of image data
3. Enable interactive consultations between radiologists and other physicians on the network

Figure 1 shows a typical PACS configuration that consists of (1) Image acquisition, (2) Communication Network, (3) Image Display and Interpretation, (4) Image Store and Retrieval, and (5) Patient Data Interface.

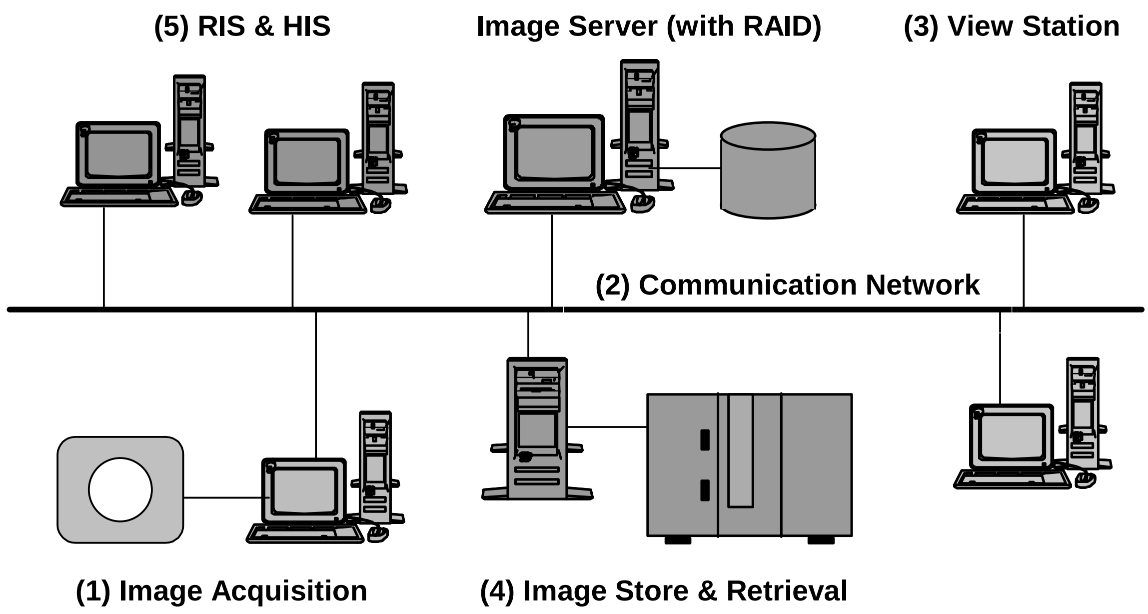

Figure 1: Typical PACS Configuration

PACS is a computer system that made up of computer hardware, software, and networks. Here are the five basic functions that a PACS performs:

1. Image Acquisition – Digital image acquisition requires interfacing the PACS to the digital imaging modalities such as CT, MRI, CR, DR, and film digitizer. The modality interfaces require that the devices to be used comply with the digital imaging and communication in medicine (DICOM) standard. DICOM represents an international standard of definitions and methods for transferring information and images between devices thus to promote reliable interoperability among different modules. For a digital modality that does not support the DICOM standard, an image gateway to convert the image is required.

2. Communication Network – Underlying the PACS is the digital communication network for transmission of images and image-related data. The structure of the network has a fundamental impact on speed of the local workstations to display new and archived images. The network function of PACS requires both a local area network (LAN) and a wide-area network (WAN). LANs enabled multiple users in a relatively small geographical area to exchange, share, and transmit both image and data on the network. WANs interconnected LANs across geographically dispersed areas such as to remote clinics and hospitals, academic medical centers and research centers. Data rates on switched LANs are approaching one gigabit per second for gigabit Ethernet. WAN's data rates range from 128K bps for ISDN to 1.5M bps for T1 line to 155M bps for ATM.

3. Image Display and Interpretation – Viewing stations or interactive workstations are connected throughout the PACS' network. Image fidelity of such a display is measured by the physical characteristics of luminance levels, dynamic range,

distortion, resolution and noise. Today, PACS display stations are broken into three types based on functionality: 1) high-resolution for primary diagnostic displays (2K x 2.5K); 2) medium-resolution for clinical review displays (1K x 1.6K); and 3) low-resolution for review displays (512 x 512). High-resolution workstations are for X-rays and replacing the conventional view box. Medium-resolution workstations are most appropriate to display CT and MRI images because those modalities acquire series of smaller images. Interactive functions allow user to perform various image manipulation and image processing with user-friendly graphic user interface (GUI). Standard interactive functions included are: contrast, zoom and pan, scrolling, image orientation, magnification, gray-scale inversion, pertinent data display, and a patient selection list. More sophisticated functions included in package like MIPAV (Medical Image Processing, Analysis and Visualization) are: image registration, segmentation, filtering and transformation, etc.

4. Image Store and Retrieval –The PACS archive should includes a long-term storage, short-term storage, and an image management database that makes use of hierarchical storage management (HSM). Most PACS employ a hierarchical storage strategy where new images are stored on a local archive of on-line hard disk and older images on a distant archive of off-line tape library for long-term storage; and the HSM manages the migration of stored images and data between the high-speed hard disk subsystem for short-term storage and the lower-speed tape subsystem for long-term storage. The selection of archiving technology depends upon archiving media and architectures. Archiving media may be in the form of magnetic discs, a redundant array of inexpensive discs (RAID), optical discs, or magnetic tape. Digital linear tape (DLT) is the favored medium due to its affordability and high reliability. Archiving architectures could be local, centralized, or distributed. The direction of archiving is toward centralized archiving, DICOM and health level 7 (HL7) interfaces, DLT tape media, and SQL databases. Other issues in archiving technology include availability, scalability, security, backup, fault tolerance, recovery, and knowledge discovery. The storage system that is functional today will be obsolete within a few years. Therefore, storage system must be able to migrate to future storage technologies easily.

5. Patient Data Interface – The hospital information system (HIS) database and the radiology information system (RIS) database must be able to interface into the PACS' image management infrastructure. HL7 is the standard for electronic data exchange of medical textual data. It was founded to develop standards for the electronic interchange of clinical, financial, and administrative information among independent healthcare computer system.

Application Services Provider

The growing role of PACS has greatly impacted the delivery of healthcare services to medical communities. However, the investment of PACS can easily run in excess of millions, and for this and other reasons many hospitals and radiological imaging centers are not quite ready to convert to the full film-less PACS. For many of these institutions and facilities, image vendors are providing the ASP as an alternate solution to lower the entrance fee into the expensive stand-alone PACS modal. ASP provides web-enabled enterprise diagnostic images distribution and archiving solution to the healthcare professionals as a service to eliminate their network and storage infrastructure and day-to-day management of the system. The applications are delivered over networks on a subscription basis. This delivery model speeds implementation, minimizes the expenses and risks incurred across the application life cycle, and overcomes the chronic shortage of qualified technical personnel available in-house. Other significant benefits that ASP provided include: a single point of support and service, lower capital investment, reduce in-house IT expertise, protection from hardware and software obsolescence, faster implementation and scalability.

There are also private practice groups feel that PACS may be well suited to one large unified institution, or several institutions under one management structure, but not for those who are working with different institutions and imaging facilities under multiple management structures. ASP model is a better fit for them since it provides all the functionality of PACS but without trying to interconnecting numerous independent hospitals and imaging centers together.

There are also hospitals that are employing a PACS system but also deploying an ASP model. They have found that ASP is a great way to archive patient medical information. They are not only saving money by not purchasing more storage devices but they are saving money on IT personnel and resources as well. They do this by keeping a large storage space on-line with terabytes of RAID, which is sufficient to keep all the most recent or short-term cases online. Previous and older studies are archived to the off-site ASP facilities over the Internet. The off-site repository is transparent and works as if it is a local repository. With T1 line, a 50 Mbytes of study may take approximately five minutes to download. With faster ATM, this could bring the delay time down to one to two minutes. Many ASP vendors (GE Medical System, Shared Medical System/Siemens, eMed and Emageon) are providing such PACS ASP solution to its customer, in which they have the diagnostic workstations and temporary storage on-site, and permanent archive and web server off-site in their facilities.

As the needs of the customer grow in the future, ASP can easily be expanded to grow with them, storing millions of images per year, without equipment obsolescence or costly upgrade paths. Client pays based solely on actual usage of the system on a one-time fee per study and usually this is less than the cost to generate film. The service includes all hardware and software needed to deploy the enterprise-wide image and data repository. Many of these ASP models come with a unique benefit of having the software and hardware reside on the client's site. Having the system on-site has advantage that the management of image and data can be restricted and complied with institute policy and procedure.

Emageon, an ASP for archive management and distribution, announced recently that it was selected as the exclusive image archive provider for the NY-based Healthcare Association of New York State (HANYS). The installation is to include more than 550 healthcare facilities in New York. Its members will have the ability to access diagnostic images securely across the Internet, intranet or virtual private network. GE Medical Systems also selected Emageon's enterprise archive software to manage the large-scale storage of diagnostic images for GE's ASP business.

In the new world of e-Health, ASP model is a powerful tool that adds flexibility and scalability to the craving healthcare community. ASP is still in the early stages of deployment and adoption. Several technical challenges remain before it will become practical. For example, questions of what, when and how to get data migrated from one physical media format to another is still a headache and requires substantial resources.

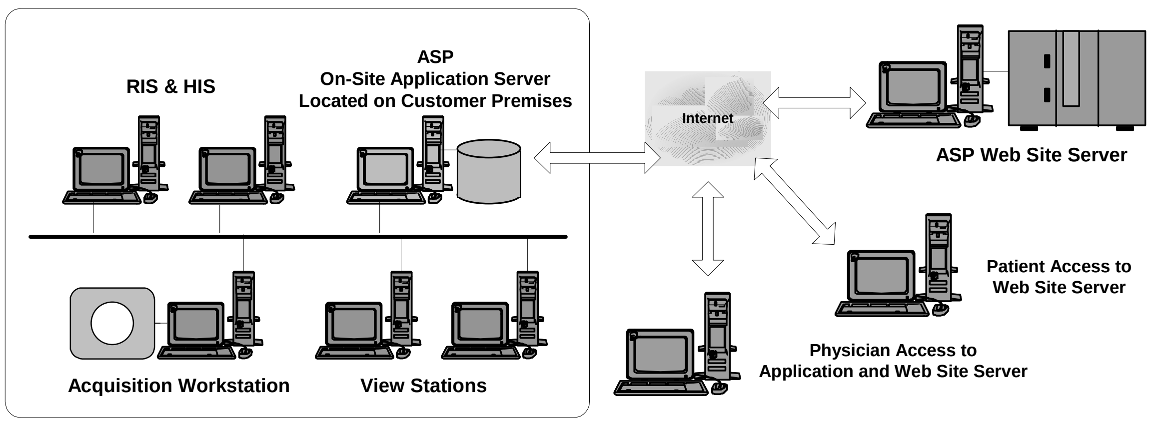

Figure 2: A Typical Application Server Provider Configuration

Medical Image Archive System

Data is a priceless corporate asset that organization needs to be captured, managed indexed, inter-related and protected. As medical image data grows exponentially, the need of a high-performance, fully centralized, fully automated and highly available storage system is required to reduce the complexity of the organization's data storage, management and retrieval. Users are also looking for a storage solution that provides faster response times and increased data availability. By integrating the PACS components and the ASP model concepts, it is possible to create an architecture that supports a center-wide archival repository.

The medical image archive system will store active (short-term) and static (long-term) data on a combination of high and low cost media in order to achieve a balance between fast delivery and low overall cost of storage. It is expected that initially the total volume of archived data will be of the order of tens of terabytes (TB), and will grow eventually to petabyte (PB).

One of the key elements of any application design is the system architecture. It describes how and what subsystems will response and interact with each other. The Medical Image Archive System should be a three-tier architecture design.

This simple architecture is illustrated in Figure 3.

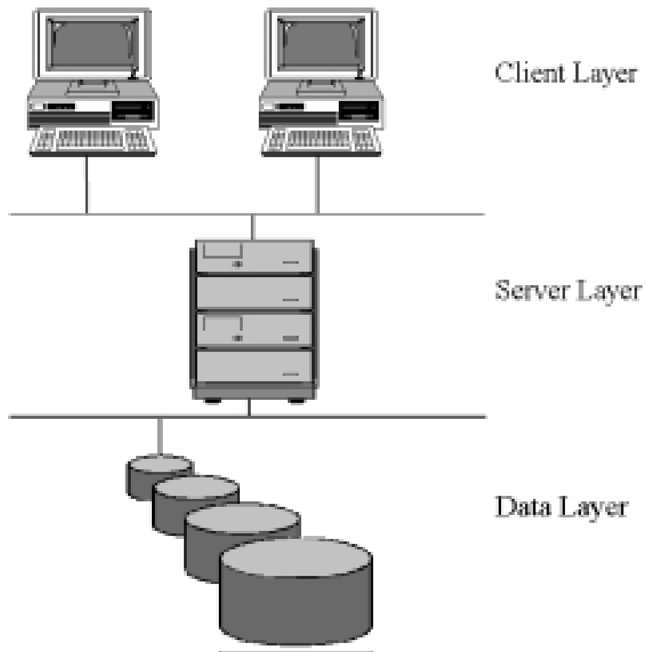

Figure 3: A Three-Tier System Architecture

The lowest tier, or the data repository layer, interacts with data usually stored in a database or in permanent storage medium. It should be able to provide automatic data management of archival and retrieval (backup, restore, HSM, etc.), and immediate and transparent data access.

The middle tier, the application server layer or the business logic layer, consists of business and data rules to assist in resource sharing. Business rules are executed on the data storage system to manipulate the database. It should be able to manage the entire workflow process from data acquisition to data repository and to automate the loading, archiving and subsequent display and delivery of medical image data.

The top tier, the client layer or the user interface layer, gives a user access to the application. This layer presents patient data and images to the user for data manipulation and data entry. The two main types of user interface for this layer are the traditional application and the Web-base application. It should be able to provide versatile data retrieval with a comprehensive managed index of archival data and a graphical user interface for seamless data access, manipulation and display.

The three tiers architecture leads itself to a three distinct phases that correspond to the different layers respectively:

(1) Identification of the storage hardware platform and a hierarchical storage management system

(2) Development of an advanced relational, object-oriented, or object-relational medical database system

(3) Building an intelligent user interface layer that allow users to capture, manage, index, interrelate, and manipulate any medical image.

Hardware Design Considerations

Network Infrastructure

The network should provide sufficient communication bandwidth that is needed to support the distribution of multimedia databases, the graphical user interface for query and retrieval of database, and the delivery of large medical images and graphics.

A typical Ethernet adapter communicates at 10 Mbits per second (10 Base T). Various other communication transmission methods are commonly used today:

(1) T1 is typically rated at 1.544 Mbps. Although, this speed is not absolutely guaranteed, it is considered very reliable. But the cost is high for most organization.

(2) Integrated Services Digital Network (ISDN) offers transfer rates of 64 Kbps. However, most telephone companies offer two channels at 64 Kbps. By combining the two channels a total of 128 Kbps is achievable. It is considered fairly reliable, but speed is slow compared to other new technologies.

(3) Digital Subscriber Line (DSL) supports a theoretical data rate of up to 20 Mbps. But the actual transfer rate usually varies depending on the quality of the wire, the distance from the phone company switching station, the computer operating system used and the brand of modem.

(4) Frame Relay is a packet-switching protocol that connects devices on a Wide Area Network, WAN. It supports data transfer rates from 56 Kbps to 45 Mbps (T3). It also allows the connection to be contained on a private network, as opposed to data being sent over a public medium such as the Internet. It is considered as reliable as T1, usually cheaper and allows for a great amount of scalability of transmission speeds.

Currently, the Institute is running on 100 Mbits per second Ethernet (12.5 Mbytes per second). For a typical medical image of 50 Mbytes, it will take approximately four second for transmission plus overhead or approximately one minute to retrieve a 500 Mbytes study. The same study delivered over the 56 Kbps modem will take approximately 20 hours.

Table 1: Mode of Data Transmission

Mode	Rate (bits / sec)	Duration
Gigabit Ethernet	1,000M	4 seconds
ATM	155 M	26 seconds
Fast Ethernet	100 M	40 seconds
ADSL	2.0 M	33.33 minutes
T1	1.5 M	44.44 minutes
ISDN	128 K	8.68 hours
Modem	56 K	19.84 hours

Storage Infrastructure

The architectural design of enterprise storage system has impacted by the fast growing of networking technology. There are three existing methods for connecting storage to computing platform: Direct Attached Storage (DAS), Network Attached Storage (NAS), and Storage Area Networks (SAN).

DAS is computer storage devises such as disk drives, disk arrays and RAID systems that are directly attached to a client computer through various adapters with standardized software protocols such as SCSI and others. This is the conventional method for connecting storage to a server. This architecture is not a networking topology; it was not designed to link multiple hosts on a common data path.

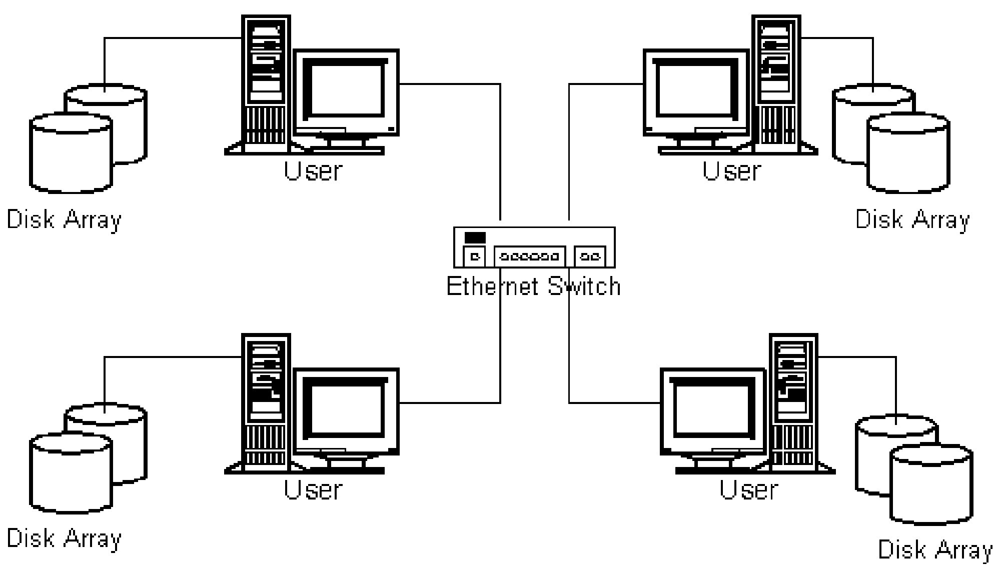

Figure 4: Direct Attached Storage Network Configuration

While the data transfer rate for DAS interconnection has gone from 20 Mbytes for SCSI-2 to 100Mbytes for Fibre Channel (FC); the data transfer rate for network interconnection has increased from a typical 12.5 Mbytes per second for 100 Base T Ethernet to 128

Mbytes per second for Gigabit Ethernet. Therefore, the bottlenecked data transfer rate has shifted from the network to the server and its direct attached storage (DAS).

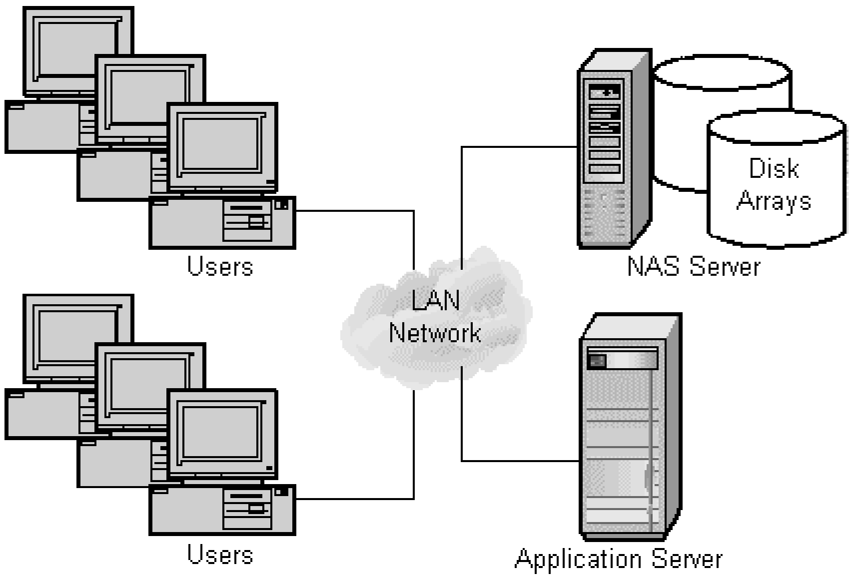

Figure 5: Network Attached Storage Configuration

NAS is storage devices that serve data directly over a network. It responds to file I/O requests coming across the network from the clients. A file system is located on the NAS and managed the device to provide a dedicated LAN-attached storage accessible to a variety of platform through standardized software protocols such as Network File System (NFS) for Unix and Common Internet File System (CIFS) for Windows. This intelligence on the NAS device enables true data sharing among heterogeneous network clients. NAS systems were designed to perform one task very well—serve files. The NAS device sits on the network and is optimized for a single purpose to serve data to users efficiently without the overhead and complexity of general-purpose servers. The offloading of storage and I/O functions to the NAS system frees the general-purpose CPU or server cycles by consolidating storage devices into smaller networks. NAS uses services that support data at the file level as opposed to the block I/O level (e.g. backup), along with sharing files across multiple computing platforms.

Instead of putting the storage devices directly on the network, SAN puts a network in between the storage subsystem and the server. This provides any-to-any connectivity for the resources on the SAN. Any server can potentially talk to any storage device. SAN boast higher availability, faster performance, centralized management, and by their design, the capability to remove bandwidth-intensive data backup and recovery operations from the Local Area Network (LAN). This frees up the LAN for normal data communications and also ensures smoother backup operations.

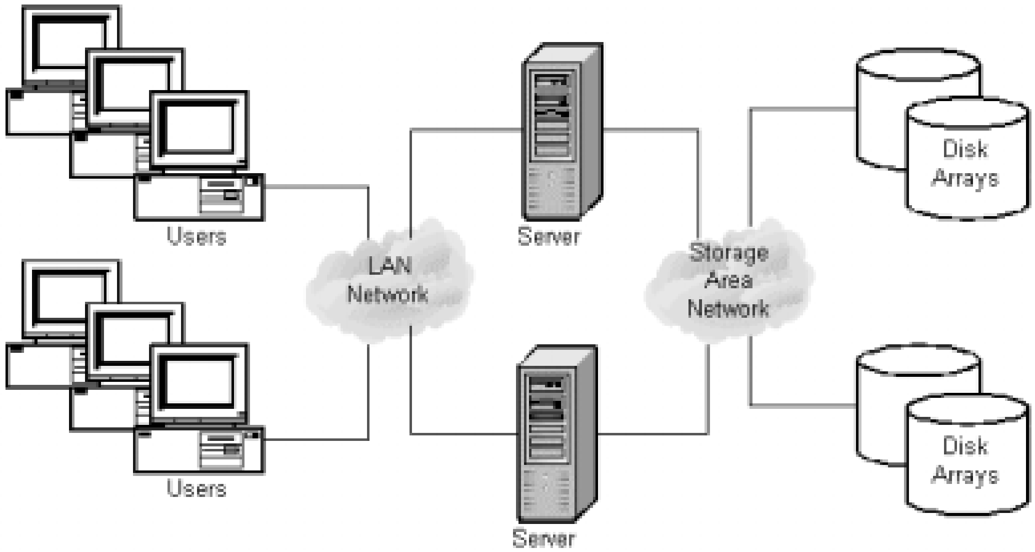

Figure 6: Storage Area Network Configuration

While the LANs may do a good job of supporting user access to servers, they are less ideal for providing servers with access to storage system. LANs are tuned for a fast user response to messaging requests and favor short or burst transmissions rather than large or continuous data transmission. Significant overhead can be imposed to move large blocks of data over the LAN, due to small packet size used by messaging protocols. Since SANs are a separate network dedicated to storage, this frees up the LAN for normal, high-performance data communications. The storage devices are not dedicated to any specific server as they are in the LAN configuration. Therefore, storage devices are external to the individual server, allowing a large-scale storage subsystem to be shared among multiple servers. The SAN architecture separates and directs network traffic to LANs and data traffic to SANs.

In a SAN, a variety of servers connect to a variety of storage devices via a Fibre Channel fabric. Fibre Channel (FC) is a network architecture as well as a transport. It provides a standard for networking, storage and data transfer. FC combines network and channel technologies to create a network for fast transfers of large volumes of information. It was developed to optimize server-to-storage data communication. It has a high-speed data transfer interface that can be used to connect workstations, mainframes, supercomputers, storage devices and displays. FC currently offers transfer rates up to 100 Mbytes per second in each direction of the fiber.

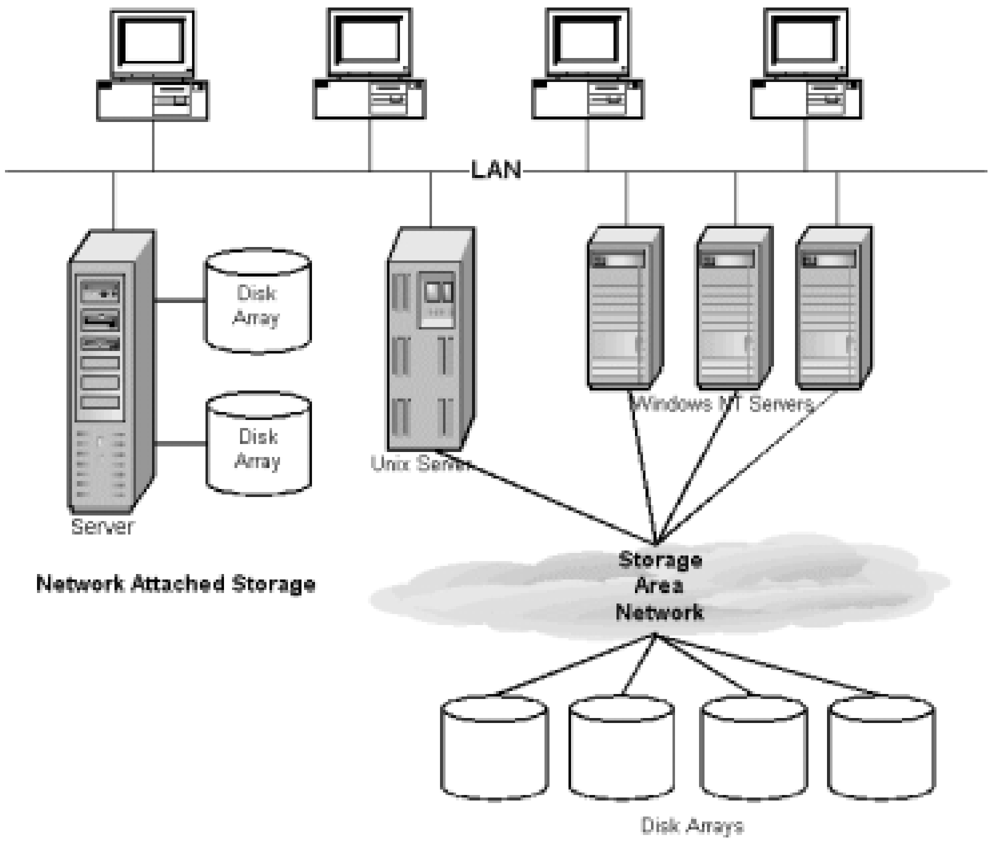

Figure 7: An Open NAS and SAN Configuration

The key element of the NAS device is an optimized server or filer. The filer is the intelligence of the device and where the file system resides and manages the actual data placement. This differs from the SAN approach where the file systems reside on each host or server and the storage volumes are attached to specific servers and the existing server operating system manages the data placement. The similarities and differences between the two technologies are described in Table 2.

Table 2: SAN vs. NAS

	SAN	NAS
Network Transport	Fibre Channel	Ethernet
Network Protocol	FCP, Serial SCSI	NFS, CIFS
Source/Target	Server/Device	Client/Server Server/Server
Transfer Objects	Block Level	File Level
Storage Device Connection	HBA	LAN Adapter
File System	The file system is located at the application server	The file system is located at the storage
Total Cost of Ownership	High	Low
Interoperability	Switches and storage devices from different	IP/Ethernet-based NAS boxes are interoperable

	vendors may not work on the same SAN	today
Multiple Operating System	SANs are volume partition-based.	Support multiple server OS simultaneously

Currently, LANs do not have the performance needed for a number of storage applications. However, as LAN performance grows to gigabits per second or 10 gigabits per second, the tendency of building a new separate network becomes less compelling.

WorldStor Inc. in Fairfax, VA., is a storage service provider, offering storage as utility to companies that have chosen to outsource their storage operations. The company spent multimillion dollars for four EMC SANs that hold between 5 and 10 terabytes of data each, EMC's Symmetrix 8730 and 8430 for storage array controllers and Connectrix Fibre Channel switches. Its Celerra File Server makes SAN volumes appear to remote clients as network-attached storages. WorldStor also makes use of EMC's Control Center management software, Symmetrix Remote Data Facility for storage mirroing, TimeFinder for replication, Volume Logix for storage virtualization and Symmetrix Data Recovery software. The high price tag is for the extensive use of EMC's software instead of developing them in-house.

On the other hand, Hannaford Bros. Co., a 112-store grocery chain purchased a Compaq StorageWorks SAN including the disk storage array, switches, host bus adapters and software for approximately $250,000. The system consists of 2.6 TB of storage, three StorageWorks RA8000 disk arrays, two eight-port Fibre Channel Switches, a 16-port Fibre Channel switch and two RAID controllers.

The intelligent of the storage is depended on the software. Storage software should be able to perform volume management, file management, clustering for high availability, backup and restore, mirroring and replication, and application management. The ability of software to make storage devices more intelligent will lead to devices that are easier to administer. Hardware vendors such as EMC and Network Appliance developed proprietary software that has the highest performance but it is tightly interwoven with their own hardware. Others software vendors like Veritas and Legato develop software-only solutions that across multiple vendor's hardware and perform nearly as well as the proprietary approaches.

Storage Devices

There are several types of electronic storage, and they are characterized by varying capacities, retrieval speeds, and costs. Today's storage material must be able to meet the high availability with faster performance, higher capacities, greater durability for data integrity, flexibility and compatibility in complex computing environment, and lowest possible cost per GB storage.

On-Line Storage

A Redundant Array of Inexpensive Disks (RAID) consists of multiple magnetic disks used in groups. RAID storage has grown to an essential part of many workstation and server installations because of its capacity, speed, reliability, and cost. RAID-based storage systems offer high storage capacity as well as good scalability for expansion. RAID systems employ a variety of data redundancy and data restoration methods for maintaining access to mission-critical data.

A majority of RAID levels involve a storage technique known as data stripping. Striping is a method of mapping data across the physical drives in an array. The data is subdivided into consecutive segments or stripes that are written sequentially across the multiple disks in the array. There are number of different RAID levels. The most common are 0, 1, 3, 5 and 0+1:

> RAID 0: Striped Disk Array without Fault Tolerance
> RAID 0 implements a striped disk array, the data is broken down into blocks and each block is written to a separate disk drive. This requires a minimum of 2 drives to implement. The I/O performance is greatly improved by spreading the I/O load across many drives in the array since read and write operation may be performed simultaneously. RIAD 0 can provide high write performance but it is not a true RAID because it is not fault-tolerant. The failure of just one drive will result in all data in an array being lost. It should never be used in mission critical environments.
>
> RAID 1: Mirroring and Duplexing
> RAID1 is the simplest form of fault-tolerant array. This array consists of multiple sets of data stored on two or more drives. The controller must perform two concurrent separate reads per mirrored pair or two duplicates writes per mirrored pair. If a drive failure occurs, data can still be retrieved from the other member of the RAID set. Mirroring is expensive since it doubles storage requirements, but it offers the ultimate in reliability.
>
> RAID 3: Parallel Transfer with Parity
> RAID 3 is same as level 0, but also reserves one dedicated disk (parity disk) for error correction. It interleaves data at the bit level across the drives in an array. It uses parity information for data recovery and stores it on a dedicated parity drive. It provides good performance and some level of fault tolerance. RAID 3 is good fit for applications that require very high data rates for a single large file because of its small strips.
>
> RAID 4: Independent Data disks with shared Parity disk
> RAID 4 is similar in some respect to RAID 3, as both levels employ two or more data disks and a dedicated parity drive. However, in this level, it uses block stripe instead of a bit level for RAID 3. If one block on a disk goes bad, the parity disk is used to recalculate the data in that block, and the block is mapped to a new

location on disk. If an entire disk fails, the parity disk prevents any data from being lost. When the failed disk is replaced, the parity disk is used to automatically recalculate its contents.

RAID 5: Independent Data disks with distributed parity blocks
This is widely used RAID. It is similar to RAID 3 except parity information distributed among all drives in the array instead of being stored on a dedicated parity drive. There must be at least three disk drives in a disk array to function as a RAID 5 disk array. RAID 5 is the most common configuration for providing drive protection at a significantly lower cost than RAID 1. The disk array controller is equipped with specialized, high-speed hardware that performs RAID parity calculations on the fly. It also features independent back-end SCSI buses that allow data to be transferred in parallel to the disk drive that comprise the redundant array. The primary disadvantage of RAID 5 is that it is not practical to add a single disk to the array because of the distributed parity. Thus if a RAID 5 implementation uses four disks in each array, then disks must be added four at a time.

RAID 0+1: High Data Transfer Performance
RAID 0+1 employs a combination of striping (RAID 0) and mirroring (RAID 1). This configuration provides high performance speed and reliability but very expensive.

StorageNetworks, a storage service provider focuses on providing storage services that model after utility companies, builds their SANs on RAID 1 and NAS on RAID 5. The charges for out sourcing range from $35.00 to $50.00 per GB per month depending on the type of service such as backup and restore for data protection and real-time data replication for zero loss.

There are other RAID configurations, but these are the types most commonly used in the industry. RAID arrays may be controlled either by dedicated hardware such as a controller-based RAID or by software such as the host-based RAID. Hardware RAID is more expensive but provides much better performance and reliability. As the cost of disk drives continues to decrease, the mirroring-based array will become the accepted redundancy technique over the parity-base array.

Medical images and data are valuable and critical to the successful operation of an organization. While RAID disk is very reliable, regularly backup of databases and applications is of important. For backup or near-line storage, a more cost effective medium is desirable, but at the expense of a time delay for retrieval. Generally speaking, less than 10% of the data would be stored on-line and the remaining 90% would be stored near-line. For such huge storage requirements a jukebox or robotic tape library is usually employed with data stored on MOD, CD, DVD, DAT, DLT or other tape format. Although magnetic tape has the disadvantage of slow and sequential access, but it is relative cheap and has a high storage capacity.

Near-line Storage

DLT is a widely used tape drive for backup as well as for archive. DLT is designed for heavy-duty usage storage environment. It has a data transfer rate of 5 MB per second and a capacity of 35 GB. The tape lifetime rated at one million passes, a media lifetime of 30 years, a head life of 30,000 hours, and a 200,000 hour MTB.

DTF is primarily directed to the professional digital video market. It is based on a 19mm-wide tape. It has a data transfer rate of 12 MB per second and a capacity of 42 GB.

Advanced Intelligent Tape (AIT) is a 8mm technology, Sony's latest solution for providing fast access high-density tape technology. In order to backup the large amount of data that can be stored in a storage hub, a library is needed to automate the movement of tapes into the drives. AIT2 offers a capacity of 50G bytes per tape. AIT3 and AIT4 formats are under development and will offer tape capacities of 100G and 200G bytes per tape respectively. Sony also said it has succeeded in recording data at a density of 6.5G bytes per square inch on a newly developed tape which translate to a capacities of up to one terabyte of storage. With more data on fewer tape cassettes, more data allow to be remaining online for a given number of drives so that faster data access can be achieved.

For disturbed or decentralized archive, jukebox for optical disk is another possibility to provide smaller storage archive. These jukeboxes can handle MOD and WORMS and are used mainly for browsing archives. Jukebox does not have very fast transfer rate but do not need to wind and rewind as tape do.

The near-line archive can consist of various levels represented by different storage material. More than one near-line archive level can be used for different access time to different content. Table 3 summaries the speed and capacity specifications for various off-line archiving media.

Table 3: Capacity and Speed of Various Offline Data Archiving Media

Device	Capacity	Data Access Speed
DAT DDS3	12-24 GB	1 MB/s
CR – ROM – RW	640 MB	
DVD	15 GB	
Magneto-Optical	2.6-5.6 GB	
DLT	35 GB	5 MB/s
DTF	42 GB	12 MB/s
AIT 1	25 GB	3 MB/s
AIT 2	50 GB	6 MB/s
AIT 3	100 GB	12 MB/s

A benchmark for backing up 1.2 TBytes of data on a StorageTek TimberWolf 9710 with ten DLT drives and 27 tape cartridges took 7.25 hours. That yields a throughput of 168 GB per hour.

Storage Management Software

Storage software that can perform file and volume management, clustering for high availability, backup and restore, mirroring and replication for disaster recovery, and application management such as databases are critical to the success of a storage system.

SAM-FS (Storage and Archive Manager – File System) software is a high performance file system and volume manager with fully integrated storage and archive management features that manages and protects large volumes of data and an unlimited number of files.

SAM-FS automatically and transparently copies files from expensive on-line disk to less expensive automated storage media, and restores the files back on-line as needed. SAM-FS allows system administrators to make very specific policy decisions for data migration and archive decision based on pre-defined criteria such as reaching a defined water mark, logical file groupings, or a combination of file size, file name, user name, last update, etc.

SAM-FS is used in data management of medical images and records because it involves a variety of files (MRI, CT, and medical records) coming from different departments and institutes. Policy decisions in this environment may involve "pre-staging" all files associated to a patient so that all the related files are fetched together.

There are four major functions of the SAM-FS advanced storage manager.

1. Archive – provide data protection by automatically copies files from disk to archival media using a wide variety of management policies. SAM-FS copies the data from on-line storage to up to four removable media volumes and leaves the original data and its metadata also on disk storage.
2. Releasing – automatically maintains disk space at specified threshold by clearing files that have been archived. It frees up on-line disk space by removing data that has been archived. Although the data has been removed, the metadata remains on-line. To the user the data appears to remain on-line.
3. Staging – automatically copies files from archival media to disk or directly to the user application. When a released file is accessed, SAM-FS stages or restores the file to disk cache. It can even allow user to begin reading a file before it is fully restored.
4. Recycle – automatically clears expired archive images from archival media to restore space. When a storage media exceeds its high threshold, the Recycler searchers for media with small amounts of useful data, copies the useful data to other media and recycles the original medium and freeing it up for reuse.

SAM-Segment provides users with an intelligent system that will copy only the segment of a file that has changed. SAM-Segment also allow users to stripe single files across multiple tape drives for reading and writing, drastically improving read and write performance. By segment very large image files, users can search and retrieve data faster because segmented files can be read back from multiple drives simultaneously.

SAM-Remote allows remote servers to migrate data from its own network directly to the tape library attached to the main SAM-FS server. For the Iceland AGES project, a remote server can attach to the local PACS system and migrate the image data to the NIH centralized medical image repository.

Mission-critical databases are often large. To maintain such databases to be available continuously, the duration to restore any unexpected failure caused by either hardware or software must be kept to minimum. Software that supports snapshot and multiple backup devices could minimize the effect on users in case of system failure.

The high availability of any archive system lies not only in the storage subsystem, disks and controllers but the host server as well. Server failover clustering allows one failover cluster to fail over to any other node in the failover cluster configuration. In this way another server can access the data immediately. This will minimize system down and provide high server availability.

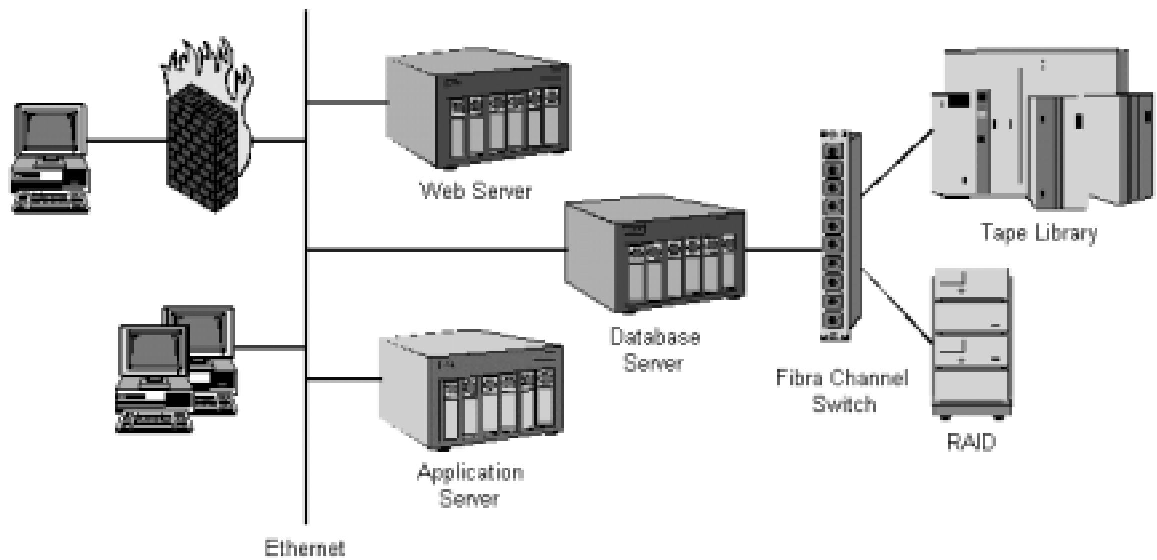

Figure 8: A Conceptual Medical Image Archive Diagram

Software Design Considerations

The software architecture defines how pieces of applications interact with each other, and what functionality each piece is responsible for performing. The application can consist of four layers, the client layer, the application layer, the database or repository layer and the archive management layer.

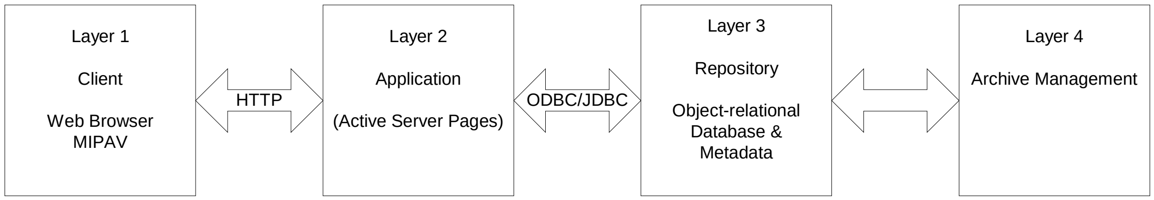

Figure 9: A Four-Tier Software Architecture

The client can be a graphical web browser that supports HTTP protocol with a DICOM front-end interface or other software such as MIPAV (Medical Image Processing, Analysis, & Visualization) that acts as the front end to the system. The application layer, which consists of web server or application server, responds to the requests submitted by clients by interfacing with the database system and applying logic to the returned results. The database layer can be a relational database or object-relational database such as Oracle that contains the medical images and meta-data and responds to the request made by the web or application server. The archive management layer performs all the behind the scene maintenance task such as backup and restore.

For example, end user on the client layer can access the medical image metadata on the web browser with Active Server Pages (ASP) procedures stored on the application layer. These ASP scripts can dynamically construct HTML documents from information returned from the database layer to the web browser client.

Java-Applets can also be used on the client layer with Java Servlets or Java Applications on the web server or the application layer and the Oracle database on the repository layer. Each application can use a JDBC (Java Database Connectivity) for executing SQL statement and database connection.

Database

Relational databases have been used to store a variety of “structured” textual information and numeric values for business applications. However, with the advent of Internet and web, the ever-increasing range of “unstructured” multimedia formats such as images and audio and video clips, require the database to have the ability to load, store, and access all the multimedia data in addition to the traditional structured data.

Database such as Oracle now supports large objects (LOBs), a multimedia file types, which can integrate unstructured data type with structured data. BFILE, a data type of LOB can be stored outside the database on CD-ROM, PhotoCD, or directly in the file system of RAID as a flat file. The advantage of storing just the locator or pointer to the image in the database rather than image itself can save a lot of disk space if an application contains many rows that point to the same digital image. It also allows other servers to access the data without going through the database. The disadvantage of storing image as a flat file is that it lacks the capability of searching the data repository by content. BLOB is used for storing the image data within the database under transactional control. The LOB can store data or image up to four gigabytes in size. Whether image data is stored within the database or outside the database, BFILE image data can still be imported into BLOB if desired.

As the size of image database grows, browsing a large number of images for a specific image becomes impractical and inefficient. Effective image database techniques attempt to reduce the search population to a feasible size and to present the user with the most promising target images.

The most common technique used is to assign descriptive metadata to index the image database in the form of keywords. Metadata refers to the descriptive information or attributes about the image data. These descriptors or attributes form the main index for image retrieval at search time so that users can find relevant information easily and efficiently without analyzing the entire database. Indexes are usually constructed manually by analyzing textual data of a patient and visual attributes of related images as to find and organize important concepts for textual and visual queries. For a DICOM image file, the format header contains text information about the scan such as patient name, medical record number, image size, imaging modality and other demographic information. A simple index for the medical images can be constructed with the DICOM header.

Two steps that are import for image database is image identification and image query. Identification of images can be broken into three categories or levels of abstraction:

- Primitive (content dependent): features such as color or shape.
- Logical (content descriptive): features that identity of the object.
- Abstract (content independent): attributes that depict the significance of the object.

These identifications are important for retrieving images on the basis of image features which current indexing practice relies heavily on.

Query of images can also be broken into five categories:

- Retrieval by browsing: provide users a thumbnail view of images.
- Retrieval by object attributes: match images that have the same attributes.
- Retrieval by shape similarity: match images that have the similar shapes
- Retrieval by spatial constraints: consider the spatial relationship between objects within an image, such as overlapping.
- Retrieval by semantic attributes: match based on the user's perception and understanding about the image.

Other techniques such as content-based image retrieval (CBIR), or query by example offer investigators the ability to retrieve images with similar etiology, but they are typically based on automatically derived features such as color, texture, shape or a combination of these. While CBIR systems currently operate effectively only at this lowest level. It suffers a significant limitation in the efficiency of retrieving the set of stored images most similar to a given query at a higher level.

For CBIR implementation, a system can extract text information about the image from the DICOM header, while another part of the system can analysis the color, texture and shape of the image or perform other image processing techniques on the image. All the result is then saved into the database.

CBIR is advancing in research and development with considerable potential. It should be exploited where appropriate. It can also be used to enhance the effectiveness of the index-based text and image retrieval systems.

Oracle8 VIR (Visual Information Retrieval) tool has some low level support for content-based retrieval based on intrinsic visual attributes of the image. It does this by identifying key features of the image such as color, texture, and structure (shape).

DICOM

DICOM (Digital Image Communication in Medicine) is the industry standard for the transmission of digital images and other associated information between digital imaging equipment. DICOM enables digital communication between image acquisition devices, image archive, hardcopy devices and other diagnostic and therapeutic equipment that is connected into a common information infrastructure. It defines data structures (formats) for medical images and related data, network-oriented service such as image transmission and query of an image archive, formats for storage media exchange, etc. DICOM file contains both a header, which include text information such as patient's name, modality, image size, etc., and image data in the same file. The DICOM standard has become the predominant standard for the integration of digital imaging systems in medicine.

Compression

The objective of image compression is to reduce the data volume by eliminating the redundancies within the image and without perceived loss of image quality. This in term allows the image to be transmitted across the networks more rapidly than its uncompressed version.

Lossless compression is important to diagnostic quality images since compressed-reconstructed image matches the original, and lossy compression is generally avoided. An ideal medical image compressor should also be able to compress images up to 16-bit grayscale and 48-bit RGB images, at a relatively high compression ratio, and with a high-speed compression and decompression rate.

Other compression options that enables users to designate different levels of compression for different imaging modalities and progressive decompression for online interactive browsing that let compressed image data decode quickly from coarse to fine or at users selectable rate would enhance the flexibility of the system.

Security

Security is critical in protecting medical information whether in paper form or electronic form. The use of password over network may not be adequate to protect the sensitive nature of medical information.

The proposed rules and standards for security and privacy of individual health record of the Health Insurance Portability and Accountability Act (HIPAA), circulated by the U.S. Department of Health and Human Services (HHS) is to protect patient healthcare information and called for stiff fine and imprisonment for wrongful discloser of information. Regarding this security standard, the question is whether or not medical images are part of the patient's medical record. HIPAA covers data security and patient privacy in three main ways: Procedurally (e.g. audit trails and security plans), Physically (e.g. disaster recovery, data backup, and access controls), and Technically (e.g. authentication, and encryption).

The ability to exchange data electronically securely over the network is essential for any healthcare organization. HIPAA compliance requires a secure solution such as Public Key Infrastructure (PKI) or digital certificate for identity authentication and Secure Sockets Layer (SSL) protocol and 128-bit encryption or Virtual Private Network (VPN) for data encryption during transmission across the Internet. Since the majority of browsers do support SSL, everyone can benefit from the use of SSL encryption.

Conclusion

The need for a medical image repository is growing at a rapid pace as more and more healthcare professionals utilize imaging for diagnosis and research. Although it is a complex task to build an enterprise class system that is reliable and robust, but the benefits it would bring to the research community is unimaginable. The following but a few examples of how other companies had done it.

Microsoft, in collaboration with Compaq, United Stated Geological Survey (USGS), and SPIN2, created a TerraServer that provides satellite imagery over the Internet. The TerraServer is a multimedia database that stores aerial and satellite images of the earth. It claims to be the world's largest atlas, containing 5TBytes of uncompressed image data. The system run on Microsoft Windows 2000 Datacenter Edition operating system, using Internet Information Server for the web server and SQL Server 2000 for the database server. On the hardware side, the database software runs on the Compaq ProLiant 8500 processor (each processor has eight 700 Mhz processors and 4GB RAM). The web server runs on the Compaq DL360 processor (each processor has two 700 Mhz processors). The storage system runs on the Compaq StorageWorks Enterprise Storage Array (RAID 5), and StorageTek 9710 tape library for backup.

Bristol-Myers Squibb is a worldwide health and personal care company. Before they centralized their data repository, they had a wide range of hardware such as SGI, Sun, and Cray computers, and they were geographically distributed throughout their facilities. Now, they have created a central repository at three of the Pharmaceutical Research Institutes for all their mission critical documents, files, and databases that could be used by their scientists to bring drugs rapidly to market. Each repository consists of a 400-500 GB of MetaStor RAID disk array, running on a dual processors Sun Ultra Enterprise 450 Sparc server and an archival storage managed by robotic DLT tape drives from Mountaingate and ATL along with a HP 320 SureStore optical jukeboxes that provide 10-12 TB of backup and archival storage. It also uses LSC's SAM-FS Hierarchical Storage Management software.

The Center for Advanced Medical Technology at the University of Utah School of Medicine in Salt Lake City initially used an optical disk system for image storage, but found that it was insufficient both in term of capacity and fast access. Currently, their system included a Sun Enterprise 5000 server with two terabytes of RAID 5 disk array for online access, a 20 terabyte of StorageTek TimberWolf DLT tape library that uses six DLT 7000 drives for long-term storage, and LSC's SAM-FS software for backup management The hospital has a ATM network backbone, and has planned to upgrade it to a full Gigabit capacity.

References

1. Amit Mehta, Keith J. Dryer, “Web Based Images Distribution.”
2. Auspex, “A Storage Architecture Guide.”
3. Bernie Hurley, “Interface and Repository Considerations for a System Architecture to Support a Distributed Digital Library.”
4. Brad Kline, “Distributed File System for Storage Area Networks.”
5. Dave Fetters, “Building a Storage Area Network.”
6. David Avrin, “System Planning and Design.”
7. Douglas Page, “Ambitious PACS Project Integrates Hospital System,” September 2000.
8. F. Unglauben, W. Hillen, M. Murdfield, “Evaluation of Two- and Three-Tier Database Connection for a Java Based Medical Image Viewer.”
9. Girish T. Hagan, “Integrating PACS into the Enterprise.”
10. H. K. Huang, "PACS: Basic Principle and Application," Wiley-Liss 1999.
11. Hyungjoo Em, “Internet Digital Library System,” Department of Math and Computer Science, Kent State University, May 1999.
12. John P Eakins, Margaret E Graham, “Content-based Image Retrieval,” Institute for Image Data Research, University of Northumbria at Newcastle, January 1999.
13. Katherine P. Andriole, “Using Staged Archives.”
14. Keith J. Dreyer, “Basic Distributed Medical Imaging.”
15. Keith J. Dreyer, “Enterprise-Wide Image Distribution.”
16. Kent Angell, “LSC’s New Solaris-Based Storage Management Product.”
17. Leann C. Beird, “Bringing all Modalities Online: A PACS Success Story,” Applied Radiology, December 1999.
18. Marc Farley, "Building Storage Networks," McGraw Hill, 2000
19. Neil F. Johnson, “In Search of the Right Image: Recognition and Tracking of Images in Image Databases, Collections, and the Internet.” Center for Secure Information Systems, George Mason University, 1999.
20. Parris M. Caulk, “The Design of a Petabyte Archive and Distribution System for the NASA ECS project.”
21. Paul Chang, “Enterprise-Wide Image Distribution.”
22. Piotr J. Slomka, Trevor Cradduck, “Experiences with a Mini-PACS System for Nuclear Medicine.”
23. Samir Shah, “Leveraging Oracle8 LOBs with Java and JDBC.”
24. Samuel J. Dwyer III, “Are You Ready for PACS.”
25. Samuel J. Dwyer III, “Network wise.”
26. Samuel J. Dwyer III, “Security-wise.”

27. Shawn H. Becker, Ronald L. Arenson, "Cost and Benefits of PACS," Journal of the American Medical Informatics Association. 1994;5:1.
28. Stephen T. C. Wong, Donny A. Tjandra, "A Digital Library for Biomedical Imaging on the Internet."
29. Steven C. Horii, "A Nontechnical Introduction to DICOM."
30. Tom Barclay, Jim Gray, Don Slutz, "Microsoft TerraServer: A Spatial Data Warehouse," Microsoft Research, Advanced Technology Division, June 1999.
31. Tom Clark, "Designing Storge Area Networks," Addison-Wesley, 1999.
32. Wayne T. Dejarnette, "PACS & Technology Update: The Next Generation."
33. Wayne T. Dejarnette, "Web Technology and its Relevance to PACS and Teleradiology," Applied Radiology, August 2000.