May 2002

Evaluation Report

NUMBER
MDA 02044

A Beginner's Guide to PACS

MDA Evaluation Report
MDA 02044

£50

Lead Author

Jonathan E. Turner (PACSnet Principal Evaluator)

Contributing Authors

Dewinder S. Bhachu (Head of PACSnet)
Alan R. McBride (PACSnet Senior Evaluator)

Centre Director

Dr. Alan Britten

PACSnet
BENCE-JONES OFFICES (PERIMETER RD.)
ST. GEORGES'S HOSPITAL
LONDON
SW17 0QT
TELE: 020 8725 3315
FAX: 020 8725 3293
WEBSITE: HTTP://WWW.PACSNET.ORG.UK

Contents

This guide was written with the aim of being a basic beginner's guide to PACS in the UK National Health Service, to explain some of the concepts involved in PACS, and to aid the process of purchasing and running a PACS. It starts with an introduction to the PACS concept and a brief history of PACS. It then discusses how images are acquired onto the PACS, how these images are stored and managed, how they can be viewed, and the movement of data between the various components of a PACS.

Some limited discussion of technical details is included, since an understanding of these details allows for a more sophisticated evaluation of system requirements and of the best way to fulfil user needs.

1. Introduction: The PACS concept

What PACS does (or is supposed to do!); advantages and disadvantages

Advances in digital technologies, particularly in the fields of computing, imaging and in communication, have progressed to the point that it is now possible to acquire medical images in digital form, archive them on computer systems, and display them in diagnostic quality. The display monitor used to present the images can be at an adjacent or distant location to the original point of acquisition. Indeed, there can be multiple monitors at multiple locations, since once the 'master' image file has been archived, it is only ever a copy of the data that is transmitted for display.

There is no universally agreed definition of what a "PACS" is or does. A Picture Archiving and Communications System (PACS) typically comprises data storage devices, image display devices, database management software and links to image and/or image data acquisition devices, connected by computer networks. There may, in addition, be network connections to other information systems such as the Hospital Information System (HIS), Patient Administration System (PAS) or the Radiology Information System (RIS). For the purposes of this document, a PACS is considered to include image display devices, data storage devices, database management software, film printers and digitisers, and computer networks. Excluded from our definition of a PACS are other information systems (e.g. HIS or RIS), image data acquisition devices (excepting film digitisers), and telemedicine systems.

There are several advantages claimed for PACS. These include:

- the increased efficiency of acquisition, viewing and reporting of images
- efficient data management, including fewer lost films
- cost savings
- space savings
- environmental benefits
- increased efficiency of retrieval of historical images

Potential disadvantages include:

- high capital investment, and ongoing costs (training, specialist staff and equipment, and so on)
- the necessary infrastructure may not be available – PACS can place high demands on computer networks, and so it may be necessary to install a new network
- the technical skills required for the support of a PACS may be in short supply

Each of the above points will be considered in detail later in this document.

Although many benefits and cost savings are claimed for a PACS, on their own these may not be sufficient to justify the large capital investment required. However, PACS can be seen as one of the first steps on the road towards a fully digital hospital data management system. With the advent of the Electronic Patient Record, it may be that a PACS will become necessary for the efficient management of image data, in much the same way that RIS and PAS are required for the efficient management of patient data, and that financial considerations will become secondary.

A PACS can be designed in several ways, the most fundamental differences between the different models being the arrangement of image data storage. The image data storage can be centralised or distributed; each model has advantages and disadvantages and these are discussed in Chapter 4, "The Image Archive", and in Chapter 10, "Networks".

Several standards have proved themselves useful in implementing a PACS, notably Digital Imaging and Communications in Medicine (DICOM) and Health Level 7 (HL7). Of increasing importance are the Integrating the Healthcare Enterprise (IHE) initiative, and the Clinical Context Object Workgroup (CCOW) standard. All these are discussed in Chapter 9, "Standards and Formats."

■ PACS and workflow

In a typical radiology department, it is likely that a large number of steps are performed in the sequence of events from the time that the patient is first registered in the department to the time that the clinical report is issued. These steps may include entry of patient details, booking of examinations, notification of arrival of the patient in the department, retrieval of historical images, dictation of the report, and so on. The required sequence is likely to be a process that has evolved over many years, and it may no longer be the optimal process for a modern radiology department.

Installation of a PACS gives the opportunity to re-evaluate the workflow within the radiology department. Rather than merely mimicking an existing, paper-based, system, a carefully planned PACS implementation can encourage improved workflow, i.e. the more efficient flow of information, images and patients through the department.

■ "Mini-PACS"

A "mini-PACS" generally refers to a small-scale PACS, usually limited to one imaging modality or department. In a typical mini-PACS installation, one or more image acquisition devices of the same modality are connected to a local digital storage device, and one or two workstations are attached to the storage device, allowing retrieval and display of stored images.

A mini-PACS has some advantages and disadvantages over a full PACS. Advantages include the relatively low initial investment compared to a full PACS implementation, the possibility of a mini-PACS growing in to a full PACS, and the chance for hospital staff to gain familiarity with PACS concepts. Disadvantages are centred on the fact that the benefits of a PACS are not fully realised – for example, if images need to be viewed outside the area covered by the mini-PACS, film still needs to be printed. Furthermore, the economies of scale of a full PACS installation are not available to a mini-PACS installation.

■ Teleradiology

Teleradiology can be defined as the ability to acquire a diagnostic image at one location and review it at a remote location, that remote location not being attached to the same local area network as the original image acquisition device. The basic components of a teleradiology system are the image acquisition and sending device, the network used to transfer the image data, and the receiving and viewing station.

Although there are many parallels between PACS and teleradiology, and in the equipment used to create PACS and teleradiology systems, there is a fundamental difference in the ultimate aim of the two systems: a PACS is designed to store image data and to transmit that data for display within an institution; a teleradiology system transmits image data to a remote location for diagnosis, but may or may not include an element of archive storage.

2. A brief history of PACS

A brief run-through of a few key dates & sites.

The concept of "digital radiology" has been in existence since the 1970s, although technological limitations meant that implementation of digital archiving systems was not practicable until the early 1980s. Even then, it took some time for PACS technologies to mature and it wasn't until the early 1990s that the technical challenges involved in achieving a fully digital radiology department were fully understood and solved.

Early PACS designs were generally based around a centralised system where images and data are stored in a central archive and distributed to display workstations on demand. Later designs featured distributed archives, where images and data are stored near to where they are most likely to be required, although improvements in network capacities allow the images to be displayed anywhere on the network.

Some important dates in the evolution of PACS include:

1968 Boston airport walk-in clinic video link to Massachusetts General Hospital, Boston

1979 first digital data link between a CT scanner and a radiation treatment planning computer – Loma Linda University Medical Centre, California

1982 First International Conference on PACS, California

1983 First EuroPACS conference

1983 US Army PACS research project

1983 Fuji introduce the first commercial computed radiography reader

1985 ACR-NEMA standard published

1985 US Army PACS installations in Seattle and Baltimore

1986 UCLA, California, USA: PACS for conferences and teaching

1986 Hokkaido University, Japan: PACS implemented

1986-1989 Utrecht University Hospital, The Netherlands: PACS implemented

1987 HL7 standard introduced

1989 ACR-NEMA 2.0 standard published

1990 US Army PACS installation at Madigan Army Medical Center, Washington, USA

1990s Massachusetts General Hospital, Boston, USA: PACS implemented

1993 Hammersmith Hospital, UK: PACS implemented

1993 DICOM 3.0 (originally ACR-NEMA 3.0) standard published

1994 Baltimore VA Medical Center

1996 Hammersmith Hospital, UK: first hospital in UK to be filmless

1998 Integrating the Healthcare Enterprise (IHE) initiative established

PACS development work continues, with technological advances making implementation of a PACS simpler and cheaper. For example, in recent years data storage devices have dramatically increased in storage capacity and have also significantly fallen in price. Universal acceptance of the DICOM standard (see chapter 9) and the increasing significance of IHE and HL7 (see chapter 9) are also having a positive effect on PACS implementation. Much current development work is focussing on the areas of workflow and systems integration.

3. Image acquisition and image data storage

A brief discussion of acquisition devices: digital – CR, CT, etc; analogue – film, CRT. Analogue versus digital data. Transmission of acquired image data to storage facility.

One of the key features of a PACS is its ability to store and transfer image data. Both the storage and transfer are performed on a digital representation of the image data, and therefore all image data acquired onto the PACS has to be digital. Many imaging modalities already produce data in digital form (for example, CT or MR) and these modalities may be attached directly to the PACS (although possibly with some intermediate interfacing unit – see Chapter 11, "Interfaces to External Systems"). Other modalities generate image data as analogue film (for example, traditional x-ray) or video display (e.g. ultrasound), and this analogue data must be converted to digital form.

■ Analogue or digital?

Data exists in one of two categories – *analogue* or *digital*. Analogue data is continuously variable (think of the speedometer needle in a car, or the level of mercury in a thermometer); digital data varies in discrete, precise steps – the numbers on a digital clock, for example. Real-world data is generally analogue, whereas computer systems deal exclusively with digital data; this digital data is usually an approximation to real-world data. Better approximations depend on having finer and finer divisions in the steps between one digital value and the next; having increasingly finer values is achieved by more powerful computer storage and processing facilities.

Although the digital representation of data is generally an approximation of real-world analogue values, it has several important advantages. Digital data, once created, can be copied easily and with no loss or degradation of data compared to the original digital data file – compare this to the process of making a copy of a film, or a simple photocopy of a sheet of paper, where the quality of the copy is often significantly lower than the original.

■ Converting analogue image data to digital

- **images on film**
  Image data stored on film is converted to digital form by using a film digitiser. See chapter 5, "Film Digitisers."
- **images on video display**
  "Frame grabbers" take the analogue video signal (analogous to the signal that comes out of a conventional television aerial cable into the back of a television set) as input and convert this into digital data, which can then be transferred and stored as conventional computer data files. The need for frame grabbers is reducing over time, since many imaging modalities that create analogue video signals create those analogue signals from digital data, and allow access to that original digital data to be made.

In recent years, Computed Radiography (CR) and Direct Radiography (DR) have become increasingly important for plain x-ray imaging. Both CR and DR, together with modalities such as CT and MR, produce digital data. Other imaging modalities such as ultrasound, angiography, fluoroscopy and nuclear medicine may produce digital data or may need to have analogue data (film or video signal) digitised. Conventional x-ray imaging produces film, which will need to be digitised (see chapter 5, "Film Digitisers").

Once data is in digital form, it can be transmitted to PACS storage. This storage may be centralised, with a single, large disc-based on-line storage facility that can accept data from anywhere in the institution and similarly distribute the data, or local storage, where image data from the local modality can be stored, for subsequent distribution as necessary. Following acquisition, image data is likely to be stored on on-line, hard disk-based storage The data may subsequently be saved on near-line or off-line archive storage, which may be tape or disk based.

It should be noted that different imaging modalities produce images requiring different amounts of data for their representation. A typical chest examination generated by CR, for example, can be of the order of 8 MB, whereas a single CT slice image may only be of the order of 0.5 MB. Note that the figure for CT is for a single slice; modern multi-slice scanners can generate many hundreds of slices.

Calculation of the amount of data required for an image is straightforward: it is the size of the image in pixels multiplied by the number of bits required to store each pixel. For example, for a typical CR chest image:

Image width: 1760 pixels
Image height: 2140 pixels
Number of bits/pixel: 16 bits

Thus to store a single CR chest image requires 1760 x 2140 x 16 bits = 60262400 bits. Assuming 8 bits per byte, 1024 bytes per kB and 1024 kB per MB, this calculates to 7.2 MB per image. Lossless compression (see chapter 7, "Display Monitors and Workstations") can typically reduce the storage requirements by about a factor of 2.

In order to calculate the archive storage requirements for a PACS, the data on the typical image sizes for each modality can be used. For each modality, multiply the typical image size by the average number of images per exam, and multiply this figure by the number of examinations performed in a year. This will give the annual storage requirements. This figure can be modified according to the amount of compression to be used on the data.

The database server

The main PACS server holds a database of all patient examination information. This database includes details of patient demographics, and the examinations the patient has had since the PACS was installed (and possibly earlier examinations, if digitised). It also contains the information necessary for the PACS to be able to find the images for those examinations, and to direct copies of the images to local storage as required (depending on the architecture of the system). The server also performs routine scheduled tasks required for system maintenance, and allows the system administrator to carry out manual tasks and system tuning as required. The regularly scheduled tasks can include daily and/or weekly backups, pre-fetching of historical images for review purposes, “flushing” of the on-line storage. Manual tasks can include management of user accounts, reconciling discrepancies within the database, and managing interfaces with external information systems.

4. The image data archive

Centralised vs. distributed storage. On-line, near-line, off-line storage. Description of different storage technologies (RAID, tape, CD, DVD, MOD, etc).

An archive identifies, stores and protects data, whether this data is in electronic form or in traditional paper and film form. Whatever form the archive takes, decisions have to be made regarding the amount of storage space needed – physical storage space for archiving of traditional media or computer storage space for archiving of digital data. Thought must be given both to immediate and to future needs. In this section we will discuss only the digital archive, including a description of the calculations necessary for the estimation of digital storage requirements. There is no "correct" storage solution: most solutions are purchased on a specific installation's functional needs and cost restrictions.

One of the key features of a PACS is its ability to archive image data files and make the image data available for viewing at one or more remote viewing stations, the image data being transmitted over a computer network. The image archive may be centralised (i.e. a single, large repository for all the image data acquired from the various imaging modalities around the site) or distributed (a number of archive devices attached to the PACS network).

Medical imaging can create large amounts of image data, both in terms of the number of images generated and the size of the image files. It therefore follows that large amounts of storage are required for the image files. This storage can fall into one of three categories:

- **on-line storage.** A storage device that makes images available immediately on demand. It is generally a "RAID" (Redundant Array of Inexpensive Disks; see below) device, consisting of a large number of hard disks. Its advantage is that data can be found and delivered very quickly; its disadvantage is that it is very expensive to obtain sufficient storage to keep an entire site's archive on-line. On-line storage may also be referred to as short-term storage or working storage.

- **near-line archive.** A large-capacity storage that is capable of storing more data than the on-line storage, often several years' worth of data, which may be sufficient to store a site's entire image archive. The near-line archive is the primary archiving device for storing the master copy of acquired image data files; it delivers copies of these image data files to the on-line storage when required. Delivering data from the near-line archive to the on-line storage takes a short but appreciable period of time, generally in the order of minutes or tens of seconds. Near-line archive devices can be built using a number of different technologies, e.g. tape or disc.

- **off-line archive.** Should the near-line archive become full, a third line of storage can be used. Space is freed in the near-line archive by removing archive media that become full, and replacing them with blank media. The removed media can be stored on a shelf, and replaced in the near-line archive device should the data contained upon it be required. Note that this replacement will require manual insertion by PACS support staff, and it could be several hours before the data becomes available. The removed media should be stored in a safe place, and needs

to be managed properly so that the disk or tape containing the required data can easily be found.

Recent increases in the capacity of systems, together with price falls, are leading to a shift in thinking away from the above definitions of storage (which are focussed on the ease of accessibility of the data), to definitions based on the lifetime of the data within storage classes. As disk-based systems increase in capacity, it becomes feasible to store several months' data "on-line", with the complete archive being stored "near-line" in a tape or disk-based archive. Thus storage can be thought of as being in one of two categories: "short-term" or "long-term."

■ Storage hardware

The most common types of hardware used to provide storage solutions are those using magnetic and optical technologies. This can be further broken down into disk and tape technologies.

Disk-based storage

RAID

The Redundant Array of Inexpensive Disks (RAID) was devised at the University of California in the late 1980's as a solution to the cost of large hard disks. RAIDs are exclusively used for online storage. Instead of using a single large hard disk to provide the required storage capacity, a RAID uses a number of smaller, cheaper disks to provide, in total, the same storage capacity. A RAID is a disk system that allows online access to the data and can provide fault tolerance and redundancy through the use of data striping and data mirroring.

Data striping

This technique spreads the blocks of each file across multiple disks but provides no redundancy of the data. Each file is separated into fixed sized blocks, or stripe units, and striped across multiple disks in a "Round Robin" manner. However, the use of this striping provides no redundancy of the data. Round Robin is a scheduling algorithm in which processes are activated in a fixed cyclic order.

Data mirroring

A technique in which data is written to two duplicate disks simultaneously hence providing redundancy in the system. If one of the disk drives fails, the system can instantly switch to the other disk without any loss of data or service. The disadvantage of data mirroring is that there is no increase in speed of access with a higher cost and double the capacity being required. Disk mirroring is used commonly in on-line database systems, where it is critical that the data be accessible at all times.

RAID levels

There are number of different RAID "levels," each providing differing amounts of speed of access and fault tolerance. The most common are level 0, level 3, and level 5.

The various levels are defined as:

- Level 0: Provides data striping. This improves data transfer performance but does not deliver fault tolerance or redundancy.
- Level 1: Provides disk mirroring. This data transfer is similar to a single disk and provides redundancy in the system.
- Level 2: Bits (rather than bytes or groups of bytes) are interleaved across multiple disks. Although this has high transfer rates, it is rarely used.
- Level 3: Same as Level 0, but also reserves one dedicated disk for error correction data. It provides good performance and some level of fault tolerance.
- Level 4: Similar to Level 3, but manages disks independently rather than in unison. Not often used.
- Level 5: Provides data striping at the byte level and also stripe error correction information. This results in excellent performance and good fault tolerance.
- Level 6: Highest reliability, but not widely used. Similar to RAID 5, but does two different parity computations or the same computation on overlapping subsets of the data.

The access capability should be able to withstand disk failure and can have the facility to reconstruct the data from a failed disk.

The use of hard disks for archive storage has grown substantially, since the reduction in price of disk storage has enabled the vendors of PACS to offer a service which allows all data to be available on-line.

Digital Versatile Disk (DVD)

This is a form of storage media which uses a similar technology to CD-ROM. However it has twice the density of pits, tracks and lands on each layer of the disk and uses a more efficient recording algorithm. DVDs come in several formats: they can be dual sided, and they can have multiple layers, which in turn allows for an increase in storage capacity. However, DVD-R is only DICOM compliant when using the universal disk format (UDF) file system, with the single sided, single layer and double sided, single layer formats. DVD has the ability to store large amounts of data, up to 4.7 GB per side for single sided, single layer disks.

Magneto Optical Disk (MOD)

A type of disk that combines magnetic disk technologies with CD-ROM technologies. The MO disk uses a laser to read the disk and a laser and magnet to write to the disk using the magneto-optical effect. A double-sided disk can store up to 9.1 GB of data.

Tape storage

There are many tape formats used for nearline and offline storage:

Advanced Intelligent Tape (AIT)

AIT comes in two formats, AIT-2 and AIT-3, which can store 50 and 100 GB respectively, with a data transfer rate of 12 MB/sec for the AIT-3. The advantage of this technology over other tapes is the in-built chip in the tape housing, which holds indexing information of the data stored on the tape. This facilitates fast retrieval of data, since it is not necessary to read indexing information from the tape in order to locate the required data. This also allows for less wear and tear on the tape as limited searching through the tape is required to retrieve the data. Manufacturers of AIT claim a tape life of 30 years.

Digital Linear Tape (DLT)

A magnetic tape storage device used in the backup of large amounts of data. The tape cartridges can store from 20 to 40 GB of data with a transfer rate of 6 MB/sec. DLT tape is cheap compared to other storage media and has fast access time. There is an upgrade to the technology, referred to as super digital linear tape (SDLT), which has a capacity of 110 GB and a transfer rate of 11 MB/sec. This has full backward compatibility with DLT; therefore any investment that has been made in DLT media will still be usable in the new drives.

Linear Tape Open (LTO)

This is an open tape standard developed by IBM, Hewlett-Packard and Seagate. It initially had two formats, based on the physical size of the storage: 25 GB for the Accelis and the higher capacity Ultrium storing 100 GB. The development has now concentrated solely on the Ultrium model. The benefit of this technology is that it is a guaranteed open standard across the three companies. As part of this standard, the LTO has a built in chip similar to AIT technology. This chip has no contacts to the unit i.e. wireless system. At present the chip does not contain the functional information of the AIT, only storing cartridge ID and media error information.

Digital Audio Tape (DAT)

This is a magnetic tape storage device which was initially designed for high quality digital sound recording. It has been adopted and used for the backup of large amounts of data. DAT conforms to the DDS standard which is the physical recording format adopted as an industry standard for DAT drives. The cartridges can store from 20 GB of data, with a transfer rate of 2.4 Mbps.

Jukebox

The media used in archive storage are housed in a device known as a jukebox. An jukebox can contain one of many types of media (CD-ROM, tape or disks). The jukebox moves the media from its storage location, by means of a robot or carousel, to a reading/writing area; the time for this movement of disk to read is usually in the order of 10 – 30 seconds.

5. Film digitisers

A description of function of film digitisers – both a short technical description of the technology and the reason behind their need within a PACS.

There are a number of situations where it may be necessary to acquire images stored on film into a PACS image data store. These include the acquisition of historical images, where a patient was examined at the hospital before the introduction of digital imaging devices; acquisition of images produced during any period of failure of the PACS; and images produced at another site.

■ The technology

Film digitisers are built using one of three different technologies.

Camera

Light is shone through the film being digitised, and a digital camera captures the information. This is a low-cost solution, but it is also a low-quality solution, and digitisers built around this technology are probably not suitable for medical imaging work.

CCD (Charge-Coupled Device)

Light (usually from a specially-designed fluorescent tube) is shone through the film, and is collected by a CCD array. A CCD is a sensitive electronic device that is capable of turning light into electrical signal; in turn, this signal can be turned into digital data. The quality of the image depends on the sensitivity of the CCD array, and the size and spacing of each element within the CCD array.

Laser

A very thin beam of laser light is shone through the film; the beam is scanned across and down the film until the whole image area is covered. The light transmitted through the film is collected by a photomultiplier tube, which turns the light energy into electrical signal which is then digitised. It has the advantages over CCD of producing a sharper image, and of having a greater dynamic range, but the ultimate resolution is comparable to CCD. It is the most expensive of the three options.

Film digitisers will typically be attached to the PACS network in a similar manner to other acquisition devices, and images acquired from film digitisers onto the PACS should be managed in the same way. This will include the creation of an examination on PACS to send the image data to. One difference is that no local storage will be required, since the original film acts as its own data store. Care should be taken to ensure that when the examinations are displayed or listed on the PACS, the "examination date" is the actual date of the original examination and not the date that the examination was digitised.

Film digitisers require regular checks to ensure that they are performing adequately.

These checks should include:

- **linearity of response of output pixel value to film optical density.**

This ensures that a change in the optical density of the film from one region on the film to another produces an appropriate change in the pixel data value stored in the image data file

- **consistency of response of output pixel value to film optical density.**

This ensures that regions within a film - or on different films - that have the same optical density produce equal pixel data values in the image data file.

- **spatial resolution.**

This is a measure of the ability of the digitiser to detect small features in an image or to differentiate between two closely-spaced, but separate, features.

- **contrast resolution.**

This is a measure of the ability of the digitiser to produce a change in pixel data values for regions in the image that are of similar, but different, optical densities.

- **geometric distortion.**

This is a test of the ability of the digitiser to create digital images which retain the proportions of the original image. In particular, all straight lines in the original image should remain straight in an image reconstructed from the digital image data.

6. Image retrieval and display

Analogue film vs. digital image display. Different monitor resolutions. Workstation functionality. Use of film printers in a PACS.

■ Film

Traditionally, medical images have been acquired to film, diagnosed on film, and stored on film. Film is, therefore, the capture medium, the display medium and the archive medium. For the performance of these roles, it has advantages and disadvantages and these are outlined below:

Film as a capture medium

To capture a medical image onto film, the sequence is typically

- load an unexposed film into a cassette (ensure light-proof conditions when loading)
- expose the cassette to x-ray radiation
- develop the film by passing it through processing chemicals under controlled conditions
- allow the film to dry

Film as a display medium

Once an image has been acquired to film, display is simply a matter of allowing bright light to pass through the film towards the viewer. Typically this is achieved by placing the film on a light box. Bright light is necessary to allow contrast in dark areas of the film to be discerned.

Film as a storage medium

Once acquired, film requires little special care to ensure that the image remains readable, provided that it is stored in a clean environment. It is a very stable storage medium, with a lifetime measured in tens of years or longer. Thought does need to be given, however, to the management of the film storage – to ensure that a patient's films can be retrieved efficiently for viewing when required – and to the provision of sufficient physical space for the storage of the films. Note also that lost films are a significant problem, both films that are misfiled and those that are lent from the film store and never returned.

■ Digital imaging

There are parallels to the use of film when acquiring, diagnosing and storing images digitally.

Digital image capture

To capture a digital image using, for example, computed radiography (CR), there is a similar chain of events to those in capturing an image to film.

- load an unexposed imaging plate into a cassette
- expose the plate to x-ray radiation

- produce the image by passing the plate through a CR plate reader
- create and store a computer file containing the image data

Other imaging modalities will form the image by processes appropriate to that modality, but however the image is formed, the end product is a computer file containing the image data. The only exception to this is where computed radiography is used to print film, with the creation of an image data file being a temporary step towards production of the film.

Digital image display

Once the image data file has been created, the image can be displayed by transmitting this data file to a viewing workstation. Here the image data is rendered to form the image, and this image is displayed on the workstation monitor. Typically, display monitors are less bright than conventional light boxes, but they have image processing capabilities that allow for contrast resolution enhancement across the range of brightness. Display monitors and workstation functionality is discussed in more detail later in this section.

Digital image storage

Digital images can be stored on a number of different types of media, including computer hard disks, tapes, optical discs, compact disc. Reasonable care should be taken that the environmental conditions are appropriate for the type of storage. Generally this means avoiding extremes of temperature and humidity, and keeping the environment clean and dust-free. Management of the images is much simpler than with film, since the PACS management application will log all stored files and their location. The problem of lost films should be much reduced, since all original image data files remain permanently in PACS storage (it is only ever a copy of stored data files that is transmitted over the network for display), and there should be no manual misfiling problems.

■ Film printers

It is likely that even a "filmless" PACS will retain the ability to print to film or other hardcopy devices. The need to print film can arise under a number of conditions, including planned system down-time for maintenance or upgrades; unplanned downtime should failure of the system or a key component occur; and production of images for teaching purposes or for transfer to another site (for example, if a patient has to be moved to a hospital without the facility to acquire and view images in digital form).

Film printers require regular checks to ensure that they are performing adequately. These checks should include such tests as:

- **linearity of film optical density to pixel value.**

This ensures that a change in the pixel data value in the image data file produces an appropriate change in the optical density of the printed film.

- **consistency of film optical density to pixel value.**

This ensure that regions in the image that have the same pixel value have the same optical density when printed to film.

- **spatial resolution.**

This is a measure of the ability of the film printer to print small features in an image or to allow two closely-spaced, but separate, features in an image to be visualised separately on a film.

- **contrast resolution.**

This is a measure of the ability of the film printer to produce a change in optical density for regions in the image that are of similar, but different, pixel values.

- **geometric distortion.**

This is a test of the ability of the film printer to print films which retain the proportions of the reconstructed image. In particular, all straight lines in the original image should remain straight on the printed film.

Other hardcopy printing devices may be appropriate; these devices include inkjet printers, thermographic printers and laser printers. These devices also require regular checks to ensure that they are performing adequately, and these checks will be similar to those described above for film printers.

7. Display monitors and workstations

Requirements for PACS workstations and monitors. Quality Assurance.

For the purposes of this document, a PACS workstation is defined as a device at which the images and related textual information necessary to perform a diagnosis or for on-going treatment management are viewed. The term "display monitor" refers to the screen used to display images, and "workstation" refers to the screen, the computer and the software running on the computer controlling what is displayed on the screen. Display monitors are manufactured in a number of different sizes, orientations and screen resolutions, depending on which role they are required to perform. Similarly, a workstation may have more or less features available to the user, depending on the complexity of the tasks that the user will be expecting to carry out at the workstation. A workstation will have one or more display monitors; typically a reporting workstation will have two or four monitors, to enable the simultaneous display of multiple images.

■ Display monitors

As noted above, display monitors are manufactured in a number of different sizes, orientations and screen resolutions, depending on which role they are required to perform.

For diagnostic reporting from x-ray imaging, the highest quality display monitors are required, capable of displaying images at the highest spatial and contrast resolutions. Such monitors will generally be those referred to as "2K" monitors, i.e. monitors capable of displaying at a high resolution (typically 2048 pixels x 1536 pixels). Lower resolution monitors ("1K" monitors, typically 1024 pixels x 768 pixels) are acceptable if reporting on images from sources with a lower matrix size (e.g. 512 pixels x 512 pixels CT images). 2K monitors have the ability to display an entire image acquired from a high resolution source (e.g. a chest image acquired using computed radiography) at the original acquisition resolution. In theory, it is possible to display high resolution images at their full resolution on 1K monitors, but in this case only part of the original image can be displayed at any one time. The entire image can be viewed by dragging the image around the display, and this function is a typical feature of PACS workstations. However, it is not clear if in practice viewing high resolution images on low resolution workstations has any deleterious effect on diagnostic power due to the lack of published studies in this area.

In order to display the maximum contrast resolution, display monitors used for reporting should have a maximum brightness of at least 180cd/m^2 (American College of Radiology Guidelines For Digital Image Data Management, 1998-2001).

As well as the choice of monitor, the conditions that the workstation is used in are important. The reporting room should have adjustable lighting, so that a low level of background illumination can be achieved. If the reporting room is to be used by more than one workstation, then care must be taken that reflections from one set of screens do not show up in an opposing workstation. Reflections and other light sources that can interfere with the displayed images must be minimised by careful design of the arrangement of workstations and monitors within the reporting room.

The functionality of a workstation depends largely on the software. At a minimum, a workstation used for diagnostic reporting should have the following features:

- windowing, so that the contrast within particular regions of interest can be maximised
- magnification
- panning, so that all areas of a magnified image can be viewed
- default display protocols, to display images in a standard, most useful manner

A review workstation generally places less demands on the display monitor, and so monitors of lower resolution can be used. Typically "1K" monitors, of screen resolution 1024 pixels x 768 pixels, can be used. The workstation software generally offers similar functionality to that offered by reporting workstations.

Black and white monitors are preferred to colour monitors for display of medical images, since black and white monitors have a potentially greater spatial resolution and can be brighter than colour displays.

Recent years have seen the development of flat panel displays, which have several advantages over the more conventional Cathode Ray Tube (CRT) displays. These advantages include:

- a saving on space
- less flicker in the displayed image
- easier to avoid reflection on a flat panel display than a curved CRT display
- more stable, require less maintenance than CRT

For large CRT displays, distortion of the image can become a problem. This distortion includes barrel distortion, change in focus over the area of the display and perceived distortion due to the curvature of the glass.

It is recommended that flat panel displays are driven by digital video cards rather than analogue cards. Using digital cards will help reduce image artefacts such as noise, jitter and ghosting of displayed images.

■ Quality assurance and quality control of display monitors and workstations

It is essential to ensure that all workstations and monitors used for diagnosis and review are performing adequately for their designated tasks. The investment of staff time into performing quality assurance checks must not be underestimated.

Tests that can be carried out on display monitors include:

- **Brightness.**

Peak display brightness should be within specification (typically 180 cd/m^2 - 200 cd/m^2). The brightness should be measured at several points on the screen, and all points should be within a specified range.

- **Contrast resolution.**

A test image containing a number of steps of different brightnesses should be displayed. It should be possible to discern each step. Particular attention should be paid to the top and bottom end of the range. The SMPTE test image is a useful image to use for this test.

- **Focus.**

A visual inspection of a test image should be made to ensure that the image is adequately sharp over the display area.

- **Geometric distortion.**

A test image which includes straight lines and circles should be displayed. Ensure that straight lines are straight to within defined limits, and similarly that circles are circular to within defined limits.

- **Flicker.**

The display should be judged for unacceptable amounts of flicker. Significant amounts of flicker can result in user fatigue.

- **Dropouts.**

Visually inspect the entire screen area for artefacts. CRT tubes may suffer from small areas where dark spots are apparent, due to a lack of phosphor on the screen at this area. LCD displays can have dark or light spots where individual crystals are malfunctioning.

8. Web browser technology

Theory and use of web servers and web browsers. Advantages and disadvantages compared to a conventional PACS workstation.

As part of their product lines, many PACS vendors offer the possibility of using web browsers to view images and clinical reports. A web browser is a standard piece of software most commonly used to display pages from the World Wide Web, for example Netscape Navigator or Microsoft Internet Explorer. Use of web browsers has several advantages, both for the purchaser and for the user:

- web browsers are a simple, robust technology. They have been under development for many years, and, provided they can access the required data reliably, perform well.
- they provide a low-cost, high functionality application.
- web browsers can run on relatively low specification computers, often on computers that already exist within a healthcare institution.
- many users will already be familiar with web browsers through use of the World Wide Web, and so it is possible that little additional training will be required.
- hospital I.T. support staff are likely to have familiarity with the use and installation of web browsers.

However, although it is possible that a PACS web browser will run on a standard desktop computer, thought should be given to the monitor used to display the images. The standard monitor may not be of adequate quality for the viewing of PACS images. Another point to consider is the functionality provided by the web browser for the display and manipulation of images. The web browser is unlikely to be as sophisticated as a dedicated manipulation and viewing application.

In order for web browsers to be able to obtain the image data files required for the display of images, it is usual for a PACS that supports web browsers to include a "web server." This is a computer that is responsible for the delivery of image data to those web browsers that request it, and is frequently a machine separate from the main PACS server and image data store.

There may be concern regarding the security of using web browsers. Although standard web browsers are used, of the same type that are used to access the World Wide Web, this does not mean that everyone with access to the internet can view the data. It is possible to limit access to particular computers or areas of the hospital network. In addition, it will be necessary to type in a standard username and password before images can be displayed.

9. Standards and formats

DICOM, HL7, IHE and CCOW; DICOM-approved formats, proprietary formats, lossy vs. lossless compression, advantages and disadvantages of compression, speed of transmission of compressed vs. uncompressed images.

■ DICOM: Digital Imaging and Communications in Medicine

Typically, a modern imaging department will contain imaging equipment sourced from a variety of manufacturers, and it is likely that any installed PACS will be required to accept image data from many, if not all, of these disparate sources. Historically, transfer and interpretation of image data acquired on one system to a second system has been problematic, since there was no commonly agreed set of rules to define the format of the data files, what information the data should contain, and how the data should be transported.

To help ease the problem of data transfer between different manufacturers' systems, the DICOM (**Digital Imaging and Communications in Medicine**) standard has been established. This standard defines a set of rules which remove many of the problems involved with the transfer of data between different systems, and it is the standard adhered to by all the manufacturers selling image acquisition devices and/or PACS.

The DICOM 3.0 standard (usually just referred to as "DICOM") originated as a standard created by a joint committee of the American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA). In 1985, this committee published the ACR-NEMA Standard Publication No. 300-1985, which was revised in 1988 to become ACR-NEMA version 2.0. These early standards did not gain universal acceptance among vendors, although the standards were used as the basis for some proprietary PACS implementations. In 1993, ACR-NEMA version 3.0 was released, and at this time the standard was renamed DICOM 3.0. DICOM 3.0 has become universal within radiology. It has also been adopted in other medical fields, such as dentistry, pathology and cardiology.

A key element of the DICOM standard is the "Conformance Statement." This describes a product's precise implementation of the DICOM standard. It should allow a prospective purchaser to decide whether the product will interoperate with other existing devices, or those other devices under purchase consideration. Many PACS vendors now place their DICOM conformance statements on their websites.

The reasons for the universal acceptance of DICOM 3.0 include:

- it is an open standard, which has input from vendors, users and academia
- the DICOM standard does not define a particular architecture to be used for a modality or a PACS
- the network protocol specified runs on top of the TCP/IP protocol, the standard protocol used for internet and other networks. Thus, there is much hardware and software already commercially available to support DICOM networking
- the standard specifies the format to be used when transferring data between DICOM-compliant devices, thus enabling interoperability between devices

- the importance of the conformance statement allows manufacturers and users to check whether devices will interoperate.

A detailed description of the DICOM standard and its use is beyond the scope of this document. However there are publications that contain material useful for further reading (see references on the PACSnet website, www.pacsnet.org.uk).

HL7: Health Level 7

This is an American National Standards Institute (ANSI) accredited standard, developed to ease the problem of transfer of data between different systems within a healthcare enterprise. The standard is developed by the HL7 group, based in the USA.

HL7 aims to standardise the format and protocols used by healthcare information systems when exchanging data items. Originally developed with data transfer between such systems as HIS and RIS in mind, it has also become the main standard used for transfer of data between a HIS or a RIS and PACS.

The name HL7 is derived from the term "Health Level 7", "level 7" being the highest level of communication in the ISO-OSI seven-layer network communication model.

IHE: Integrating the Healthcare Enterprise

"Integrating the Healthcare Enterprise" (IHE) is a joint initiative of the Radiological Society of North America (RSNA) and another USA-based society, the Healthcare Information and Management Systems Society (HIMSS). It aims to develop a framework to ease the integration of, and information flow between, various medical information systems (for example, a PACS with a Radiology Information System (RIS)). Its ultimate goal is to ensure that all required clinical information is correct and readily available to the relevant users. Currently, interfacing systems can be difficult due to the different standards – and implementation of those standards – employed by such systems. IHE aims to define how existing standards may be used to facilitate communication between computer systems used in healthcare. Initially, IHE will concentrate on DICOM and HL7 but may extend to other standards if necessary.

Increasingly, computer systems are used to acquire and store patient information, test results and images. A single patient can have data stored on several disparate computer systems. To maximise its value, data from multiple sources needs to be available to clinical users at a single point of use. Standards have been developed to improve information flow within systems (e.g. DICOM within medical imaging systems, and HL7 within clinical information systems). IHE aims to improve information flow between these systems by taking existing standards and providing a framework for their implementation, so that information can be shared with maximum efficiency.

Knowledge of IHE can give many benefits to users and implementers of hospital information systems.

■ CCOW: Clinical Context Object Workgroup

CCOW is an ANSI-certified standard, published by the HL7 group. It complements the HL7 standard's focus on workflow and integration of data between systems by focusing on the presentation of data to the end-user, so that the shared data appears to have come from a single system rather than many.

The CCOW standard was originated in 1996 by a consortium of healthcare providers and vendors. The CCOW Technical Committee became part of the HL7 group in 1998 and gained ANSI certification in 1999. More information on CCOW can be obtained from the CCOW pages on the HL7 web site (links available from the PACSnet website at www.pacsnet.org.uk).

■ Image data file formats

There are a variety of file formats in use for the storage of image data. The common aim of the file formats is to save data in such a way that a representation of the original image can be constructed from the saved data. Most file formats attempt to compress the original image data in some way, in order to save on storage space and transmission times over networks. The file formats used fall into one of two categories: "lossless" and "lossy."

Using lossless compression, it is possible to re-create exactly the original data file from the saved data file. Compression ratios achievable with lossless compression are generally around 2:1 or 3:1 compared to the original file size.

Using lossy compression, an approximation of the original data file is constructed from the saved file, with the accuracy of the approximation depending on the degree of compression used. Compression ratios achievable with lossy compression vary markedly, but are generally in the range 10:1 to 100:1. The key advantage of lossy compression is that the file sizes created are much smaller than those created using lossless compression. Thus take up much less storage space and are faster to transmit over networks. A potential disadvantage is that it can take longer to decompress the data back to its original size.

10. Networks

Network topologies and architectures; protocols.

■ Introduction

The network is the most fundamental area of a PACS, as this transports the data from the servers or modalities to the displays and archives. The speed of networks has greatly improved in recent years; however the design of the network can greatly affect the functioning of the PACS. Although bandwidth (the amount of data that can be transmitted in a fixed amount of time) is increasing, it is still important to design and implement a network with much care to avoid any problems during operation.

Networks are used to link computers to computers, or computers to archives. For this to be achieved a standard has been created by the International Standards Organisation (ISO). This incorporates a model with seven layers from the top (application) layer, where the software runs to the bottom (physical) layer, which is the cable the data passes along. This is known as the Open System Interconnection (OSI) model.

In general, the seven layers are simplified to four: Application Layer, Network Protocol, Data Link, and Physical layer. This covers the application sending data using the protocol (which could be TCP/IP or Appletalk) through the network card and along the cable which is UTP (unshielded twisted pair, also known as CAT5) to the other computer. The data link layer is the network card which can be one of several types, including Ethernet, ATM, token ring or FDDI.

Data is sent over networks in blocks known as packets. Each of the packets is sent individually and may even follow different routes to its destination. Once all the packets forming a message arrive at the destination, they are recompiled into the original message. The packets vary in size depending on the network protocol employed by the network: for Ethernet, they vary in size, while for ATM networks the size is constant. Multiple devices may attached to the network, any or all of which may be attempting to send data packets at any one time. If devices do attempt to send data packets simultaneously, this can lead to “collisions” between packets, requiring data to be resent, with a consequent drop in network efficiency. As a result of this drop in efficienct, many networks claiming 10 Mb/sec are actually only functioning at 1/8 of this rate under optimum conditions.

Asynchronous Transfer Mode (ATM)

This is a network technology based on transferring data in packets of a fixed size. The size of the data packets in ATM are relatively small compared to units used with older technologies. The small, constant size allows ATM equipment to transmit video, audio and computer data over the same network.

Ethernet

Ethernet is the most widely used architecture in local area networks (LAN), and is defined by the IEEE as the 802.3 standard. Ethernet operates over various types of physical media such as coaxial, shielded or unshielded twisted pair and fibre optics. This normally operates at 10 or 100 Mbps but new versions can operate at Gigabit (1000 Mbps) transfer rates.

Equipment

A network requires specialist hardware to function. The simplest device is known as a "hub." This is a device that links a number of peripherals or workstations together through a number of ports. The hub has problems as all the packets on the network are sent to all the ports, creating a number of collisions and hence slowing data transmission.

The network can be comprised of a series of subunits, which can either be homogeneous (e.g. Ethernet to Ethernet), or heterogeneous (e.g. Ethernet to FDDI). The subunits can be connected by a bridge. A Bridge is protocol independent and does not analyse the packets sent across the network; hence it does not re-route the data and consequently it is faster than a router.

A more sophisticated way of connecting the subunits of a network is to use routers. Routers use headers and a forwarding table to determine where packets are to go, and are able to determine the best route between any two hosts. Thus, they only send data along network subunits where the data is useful, avoiding adding to the network traffic on network subunits where those particular data are not needed.

Network topology and architecture

A PACS network can be designed in a number of different ways. Three of the most common designs are described below:

Bus

This is a network topology where all the computers can access the network on a 'first come first served' basis (figure 1). In an Ethernet architecture this can lead to collisions occurring on the network when two computers send data along the line simultaneously. To compensate for this a protocol known as Carrier Sense Multiple Access/Collision Detection (CSMA/CD) is used. If two computers use the network simultaneously and this protocol senses a collision, both computers will stop transmitting and will send later.

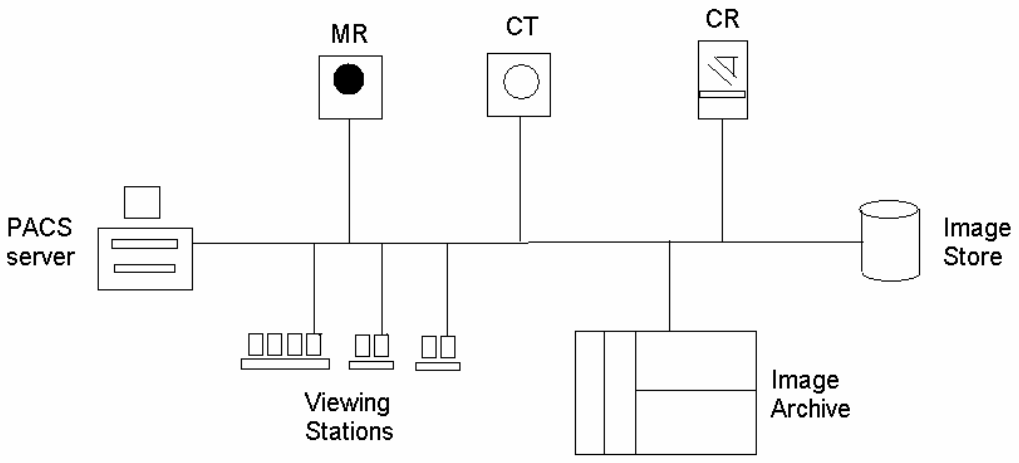

Figure 1. A PACS using a bus network arrangement.

Star

In a Star topology there is a direct connection from the computer to a central hub (figure 2). This direct connection provides what would appear to be an optimised solution that is highly reliable. However the whole system is dependent on the speed and the integrity of the central hub. The central hub can be a server which controls the network or a piece of hardware that links two networks together. The networks can either be homogeneous (e.g. Ethernet to Ethernet) or heterogeneous (e.g. Ethernet to ATM). The sub unit of the main network is referred to as a subnet.

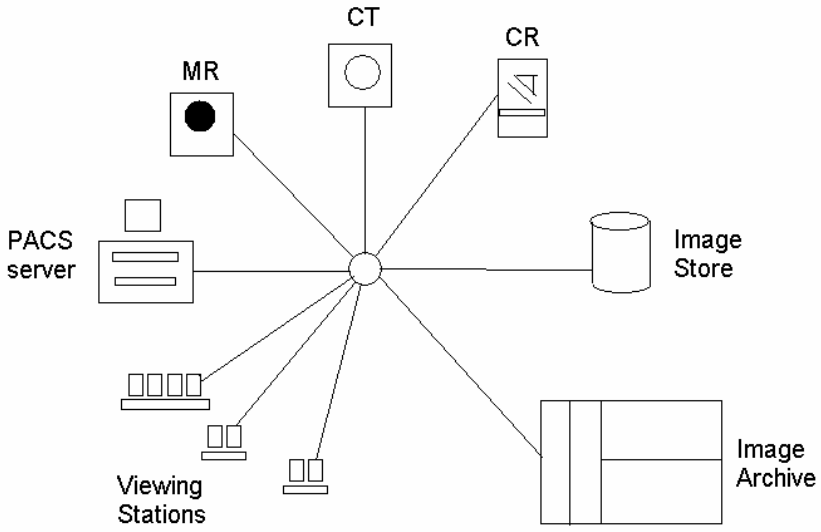

Figure 1. A PACS using a star network arrangement.

Ring

In a ring network the computers are all linked together by a single cable with the computers accessing the network on a set pattern (figure 3). The most common is the token ring, where a virtual token is passed to each computer in turn: the computer holding the token has control of the network and can send data across the network. Fibre Distributed Data Interface (FDDI) also works on a ring topology. However this has two optical fibres with the data flowing in opposite directions which allows for redundancy in the network. If one of the nodes fails it is disconnected and the other network takes control to complete the circuit.

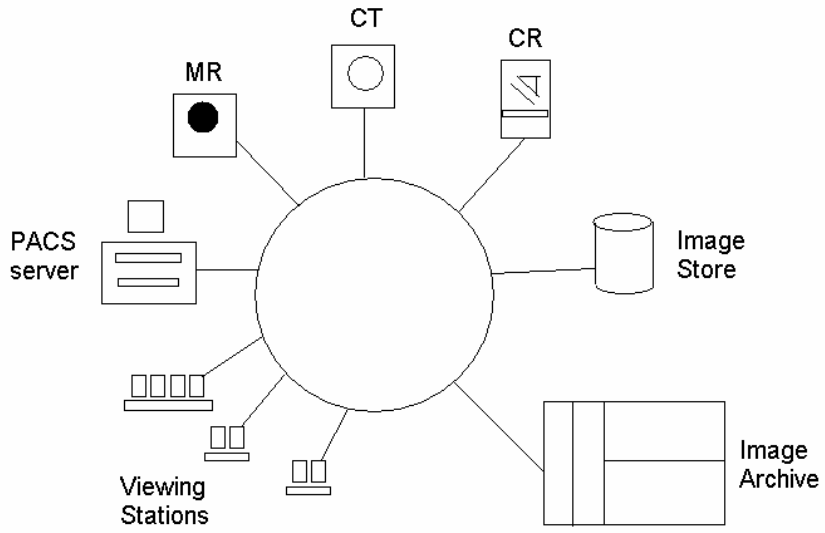

Figure 3. A PACS using a ring network arrangement.

■ Networked storage

The devices listed above are attached to the network using two differing network methodologies, Network Attached Storage (NAS) and Storage Area Network (SAN).

NAS

In NAS, the storage device is usually a RAID attached directly to the network. It is controlled by a NAS controller, which is a processor with an operating system, for example a UNIX computer. This will see all the other storage devices on the network and operates using Network File System (NFS). NFS is a file system protocol developed by SUN Microsystems™. It is now a standard protocol in UNIX that allows computers to access each other's files. Using NFS, the files on a remote workstation appear as part of the local storage.

SAN

A SAN is a network whose primary purpose is the transfer of data from computers to stage elements. A SAN provides an infrastructure which contains physical connections between the computers and the storage devices via a software management layer which organises the attached computers.

11. Interfaces to external systems

Some external systems: HIS, RIS, EPR. PACS broker. Description of the information that PACS is likely to require from external systems – e.g. patient demographics, exam orders, clinic appointments, reports.

When designing and implementing a PACS, it is important to give much thought to image data archiving needs, networking requirements, image display and management requirements. It is also important to give detailed thought to non-image data (e.g. patient demographic data, examination details, clinic appointments, etc). The storage, management and acquisition of such data are all important aspects that need to be considered.

Much, if not all, of the non-image data required by a PACS can be acquired from external systems such as a Radiology Information System (RIS) or a Hospital Information System (HIS). The functions of some external systems are described below.

■ Radiology Information System (RIS)

The core function of a Radiology Information System (RIS) is to store text information pertaining to patients and their examinations within a radiology department. Typically, a RIS will store information including patient name, date of birth, address, telephone number and information on general practitioners. It will also store information on a patient's examinations, including date and time of examination, where the exam was performed, the type of examination and the clinical report. It will also store information on future examination bookings, and may have a scheduling facility that will aid choice of booking date. Much of this information is needed by a PACS; for example, a PACS will need to know which examinations have been entered on the RIS so that an appropriate entry can be made in the PACS database.

■ Hospital Information System (HIS)

A Hospital Information System provides patient-related information to the entire hospital. Of particular relevance to PACS is the information on current patient location and forthcoming patient attendances. Patient location information is used by the PACS to provide filtered lists of examinations to the system user (for example an in-patient list, or a list of patients for a ward round). Information on forthcoming patient attendances, for example clinic visits, can be used to manage the on-line storage for those systems that have a hierarchical storage structure. If the PACS can receive information from the HIS concerning forthcoming patient attendances, the PACS can ensure that all relevant images for the visiting patient are retrieved from archive and are available on the on-line storage, thus making it possible to display the images with minimal delay.

A diagram showing typical information flows between hospital information systems is shown in figure 4. A diagram showing information flow using IHE terms (see Chapter 9, "Standards and Formats") is shown in figure 5.

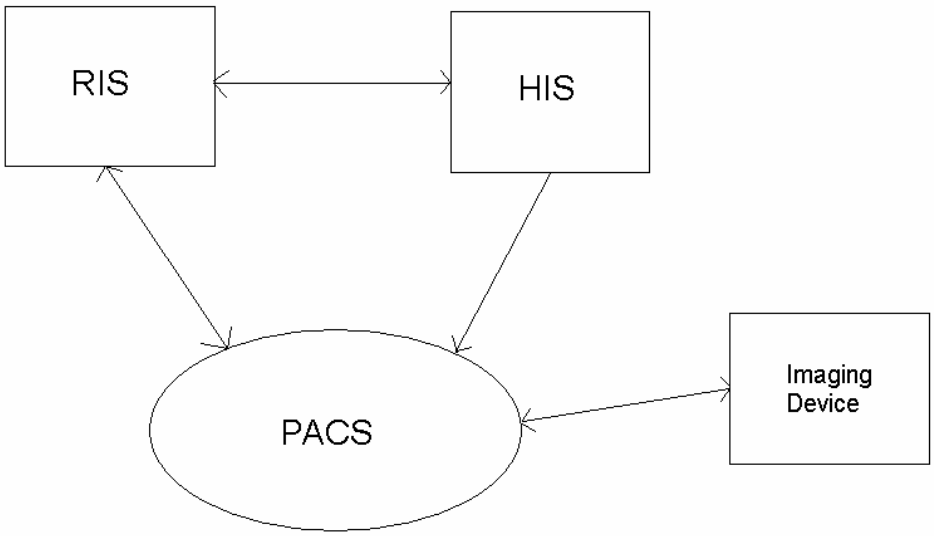

Figure 4. Information flow between various hospital systems.

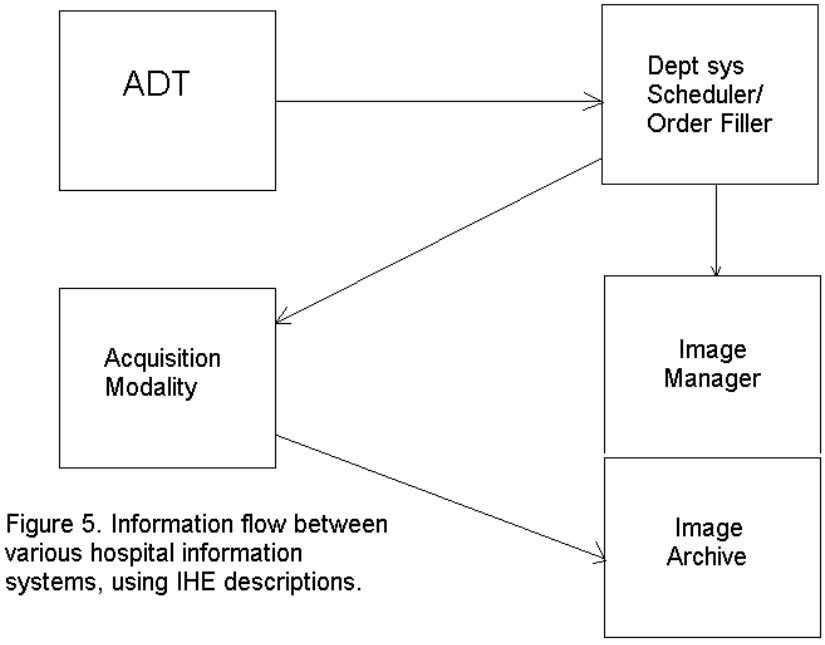

Figure 5. Information flow between various hospital information systems, using IHE descriptions.

The PACS broker

To allow for efficient communication between an external information system and a PACS, it may be necessary to use an interface engine. An interface engine is a computer system which sits between two (or more) information systems, and reformats the output from one system into a form readable by the receiving system. In a PACS implementation, this interface engine is generally referred to as a "PACS Broker". It typically receives information on patient demographics and examination bookings from a RIS and/or a HIS, and passes this data on to the PACS, following appropriate reformatting by the broker. Data can also flow from the PACS to an external system should the external system have the ability to use data from the PACS (for example information from an imaging modality regarding the acquisition status of an examination). The PACS broker may maintain its own database of information derived from information acquired from the attached devices.

Integrating other hospital information systems with a PACS has several advantages:

- a single point of entry for data means consistency, and avoids duplication of effort
- access to scheduling information aids image data storage management. Information on scheduled clinic visits or radiological examinations can be sent from the HIS or the RIS to the PACS. The PACS can then use this information on which patients will be attending the healthcare institution to ensure that all relevant images for those patients are available in on-line storage in time for their attendance.
- access to patient location information allows for filtered worklists. PACS can obtain information on patient status and location from the HIS, and use this to present useful worklists to radiologists and clinicians. For example, a list of examinations for all patients on a particular ward can be created.

The DICOM gateway

There may be occasions when imaging acquisition devices that are not DICOM-compliant will be required to send images to a PACS – for example, following the installation of a PACS in a hospital that has existing, non-DICOM, imaging devices that are not due for replacement. Under these circumstances, it can be possible to acquire images onto the PACS by use of a "DICOM gateway." In a similar manner to which a PACS broker allows information flow between a PACS and external information systems, a DICOM gateway can allow image data to flow from imaging devices to the PACS.

12. Backups

Why backups are important. Frequency of backups, backup media, where and how to store backups.

The data stored in a PACS - both image data and text-based - is very important, and is likely to be irreplaceable. Thus, it is important to ensure that the data is stored safely and securely, and that, where appropriate, backup copies of the data are made. Slightly different considerations apply depending on whether it is image data or text-based patient data that is being considered.

■ Image data

By far the largest quantity of data stored on a PACS is created by image data files. This creates particular problems when considering backups. Archive strategies have been covered earlier in this document (see chapter 4, "The Image Archive"), and the various options for archive storage have been discussed. It should be established by each site whether backups of image data are needed. In a traditional film store, there is generally only one copy of any image, and if that copy is lost then the patient must be re-examined if an appropriate image is required.

Similarly, it would be possible to design a PACS where the only permanent copy of image data is in the PACS archive. However, given the relative ease with which copies of digital data can be created and stored, it would seem prudent to investigate the options available for backing up digital image data. These include copying the data on-site and moving the media to a remote location and transfer of the data over the network to off-site backup.

Advantages of on-site storage of data backups are that they are less prone to network problems and have total control of backup management.

Advantages of off-site storage of data backups are that they can be done automatically as each image is acquired, and there is no overhead in management time.

The frequency with which backups of image data should be performed depends on several factors: the time and effort required to replace lost data against the time to retrieve it from backup; the time to perform a backup and the impact on the system performance while the backup is being carried out and the expense of performing backups, both staffing costs and the cost of the media (e.g. tapes) and storage (e.g. fire safe) required.

Consideration should be given to the impact on the PACS of any backups performed. If the system is unavailable to users while the backup is being performed, then the backup should be scheduled for a time that it will have the minimum impact on users. It may be better to perform regular, but short, incremental backups if it is considered that this will have a less deleterious impact on the overall system availability than less frequent, but longer, full backups.

■ Backup media

The choice of media for backup data is similar to the choice of media for archive storage, i.e. tape, optical disk, CD, DVD. Since the need is for storing a large amount of data, with access to any individual item of data not being a prime consideration, it is likely that tape will give the best price/performance rating of the available options. See Chapter 4, “The Image Archive.”

■ Patient data

This includes all information on patient names, dates of birth, hospital number(s), and examinations booked and performed. Since this data is textual rather than image, it takes up a significantly smaller amount of computer storage space than image data. Backups are therefore much more practicable than for image data, and in many ways the information is more critical. It is therefore recommended that regular backups of patient and associated data are taken. Similar considerations regarding the frequency of backup for text-based data apply as for image data, above.

13. Security

Security of system to intruders – both physical and via computer networks. Protection against fire etc. Access the system.

Security of data is highly important and should be given prime consideration when planning and implementing a PACS. This section discusses two separate security issues; the unauthorised access to data and the physical security of PACS components.

There are two aspects to security: hardware-based and software-based. Hardware-based security includes such tactics as housing the main PACS servers and archiving devices in a secure room with access limited to essential staff only, physically fastening workstations and monitors that are in less secure areas and concealing network cables wherever possible.

Software-based security covers important areas like user access to the data within the PACS. User access should be controlled by allocation of individual user accounts. An individual's account should allow access only to the data and system facilities that the user needs to see. Accounts should be protected by password and possibly additional physical security devices such as fingerprint identification. An audit trail, which keeps a log of all transactions made on the PACS and the user initiating each transaction, can also prove useful.

Protection of the PACS from unauthorised external access is aided by use of firewalls. A firewall is a computing device used to make a network secure. It is a combination of computer hardware and software. It works by restricting access to the network to trusted external users only.

14. The procurement process

The PaSA/PACSnet Procurement Process and the SON/DSON.

Currently within the NHS, the procurement of PACS is highly fragmented, resulting in high purchasing costs for both purchasers and providers. The main reasons for this are: a mixture of Radiology, IT and occasionally Estates led procurements, each with differing methodologies; lack of policy or procurement guidance; no standardisation of process or documentation. In order to address this, the Diagnostic Medical Equipment team within the Purchasing and Supplies Agency (PaSA) established a PACS Working Group in July 2001. The membership of the working group currently comprises:

- PaSA
- PACSnet (MDA)
- NHSE NW Regional Office
- Bromley Hospitals NHS Trust
- University College London NHS Trust
- Agfa
- Ferrania
- Fuji
- G.E Medical Systems
- Kodak
- Laser Lines
- Philips
- REAL Software
- Siemens

The working group agreed three initial (key) areas for development:

1. Procurement process
2. Standardised templates, specifications and information requirements
3. PACS-specific terms and conditions to cover capital, lease and managed service contracts

An outline of the procurement process and the technical templates has been agreed and summarised below:

At the outset of the project a "steering group" should be established with clear representation from all stakeholders at the Trust and to have a board level sponsor for the project. A PACS procurement represents a major investment. Trusts, Primary Care Trusts (PCT) and Strategic Health Authorities (SHA) will need to consider the economic impact as well as the effect on quality of care and delivery of service to patients. The NHS insures that these issues are addressed through the business case process, which has three distinct stages:

1. The Strategic Outline Case (SOC) outlining the project prior to any detailed planning

2. The Outline Business Case (OBC) defining the scope of the project, level of investment and preferred procurement route. The OBC leads to formal approval and commencement of the procurement process.

3. Full Business Case (FBC) leading to final approval upon completion of the procurement

Prior to the commencement of the procurement, the OBC must be completed and presented. The OBC must include direct cost benefits (money saved on film, chemistry, processing, storage and staffing), and indirect savings (improvements in speed, efficiency, greater patient throughput, rapid availability of images, reduction in re-takes and patient exposure). These should however, be balanced against the costs (the costs of the system, support services, installation and infrastructure requirement (networks etc) and the costs of additional staff such as a systems manager, etc).

In order to drive the procurement, a project team should be established to include as a minimum; the PACS project manager, representation from Purchasing (Trust supplies team and/or PaSA), Finance, IT and Users (Radiology, A&E, Orthopaedics, Cardiology, Theatres, Intensive Care, etc).

The procurement is initiated by the placement of the OJEU Notice (if the anticipated contract value exceeds the EU threshold of £100,410 exc. VAT). A minimum of 37 days is allowed for expressions of interest.

The process utilised by PaSA is the Negotiated Procedure since PACS contracts are solution based and as such, are purchased as the provision of services not goods. Template adverts are available from PaSA and can be submitted on a Trust's behalf. PaSA has also pre-vetted the major suppliers based on the criteria in the OJEU Notice on behalf of NHS Trusts.

Whilst the advert is being issued the project team should start to prepare the initial Statement of Need (SON). The SON provides suppliers with the basic background and information on the Trusts anticipated requirements. The SON will include details on the Trust infrastructure, objectives of the project, and linkage with other strategies such as the Electronic Patient Record (EPR) and the Local Implementation Strategy (LIS), and the DICOM status of current equipment to be connected to the PACS. To assist Trusts in establishing DICOM status, the PACS Working Group has established some generic DICOM wording, available from PACSnet or from PaSA. This can be used not only for the purchase of new equipment but to establish the DICOM status of current equipment.

Once the SON is completed and the Notice has closed, the SON should be issued to vetted suppliers. Presentations by suppliers showing their initial solutions based upon the Trusts SON should be arranged. The presentations should be used to build Trusts knowledge of the products available, the viability of their requirements and the ability of suppliers to meet them. It is as this stage that the Trust will conclude its first stage of short listing. To assist Trusts, the Technical sub group has developed evaluation sheets along with a scoring and weighting system, available from PaSA. It is essential that prior to the presentations the project team agree the weighting for each section of the Presentation Evaluation Sheet (PES) along with any specific questions. This should then be given to the suppliers prior to the presentations. The presentations should be evaluated on technical, clinical and financial elements. Although the Trust may commence site visits, it is advised that these are conducted after short listing so that only viable suppliers are visited.

It is suggested that between 3 and 5 suppliers be short-listed and site visits undertaken. Details of these visits should be evaluated and to this end the PACS Working Group has developed a site visit evaluation sheet, available from PaSA. Prior to the commencement of visits, it is essential that the project team meet and the weighting of each section of the SVES along with any specific questions be agreed. This information should be provided to the supplier prior to arranging visits to ensure viable sites are evaluated. Informal visits should also be conducted to corroborate findings.

In addition to the site visits, the SON should be developed into the Detailed Statement Of Need (DSON) based on the knowledge acquired by the Trust. The DSON builds upon the SON by updating requirements, focusing on working practices and the expected outcomes of the solution (output based specification) and a more detailed cost breakdown and exploration of financial models such as, managed services and initial contract negotiations. The individual questions included within the DSON should be amended and weighted in accordance to the Trust's requirements

This should be issued to the short listed suppliers, with the appropriate weightings, for a formal response within four to five weeks. The responses should be scored and used in conjunction with the SVES. It is suggested that a 3:1 split be applied between the weight of the DSON and the SVES. This should provide a short-list of no more than three suppliers to be issued tenders.

Both the PES and SVES provide a clear audit trail of the decision making process and should be fully completed. The template SON and DSON contain example questions/information which are highlighted in red. These should be amended to suit individual requirements.

The short listed suppliers from the previous stage (1 – 3 suppliers) should now be issued with the tender. This should include the terms and conditions and schedules to contract along with any amendment to the DSON following clarifications. The tender is essentially the formalisation of the offer(s) for which up to three weeks should be allowed (depending on the number of suppliers). The tender response should be evaluated on the basis of technical, clinical and financial grounds (the majority of the work will have been done through the preceding stages). This will allow the identification of the preferred supplier.

On identification of the preferred supplier contact negotiations can be finalised. These will be based on the Trust's DSON and the supplier's response.

The final stage in the procurement process is the completion of the full business case (FBC). Once this is approved the interconnectivity testing can be completed and the criteria for implementation and acceptance testing be established, leading to contract award.

Upon contract award the award notice must be issued within 48 days. A template is available from PaSA and can be issued on the Trust's behalf if PaSA issued the EU Notice.

15. Staffing issues

Changes in staff structure. Changes in staff tasks and responsibilities. Change to workflow.

The implementation of a PACS will have an impact both on the mix of staff supporting the radiology service and the working patterns of many members of the radiology staff.

Installation of a PACS is likely to mean the loss of an on-site film store, which will mean that staff will no longer be required to perform film filing duties. During the change-over to PACS it is likely that many films from the existing film store will be digitised, and so during this period the film store filing clerks will be needed. However, once the bulk of recent films have been digitised, it is likely that film storage and management will be moved to a remote location, possibly under the management of a third party, and so the need for film filing clerks will reduce.

A PACS is a technically complex system, and its management is a specialist task. A PACS systems administrator will be needed, and this will almost certainly be a full time appointment. Further support should also be in place for the system and its components (e.g. networking, storage, workstations, etc). This support may come from hospital IT staff, the PACS vendor or from additional staff within the radiology department.

Quality Assurance of the various components of a PACS system is also a specialist task. This work is likely to fall upon hospital scientific support staff, most probably a Medical Physicist. Quality control and quality assurance tasks that will require scientific support include checks on display monitors, and, if appropriate, film digitisers and film printers.

Other new tasks that will appear with a PACS include system administration tasks (maintenance of error logs, management of system user accounts, tidying of databases) and diagnosis of system faults (investigating any failed communications from the PACS to external systems, and diagnosis and repair of failed image transfers from an acquisition modality to PACS).

An indirect cost of implementing a PACS is training. System managers and those involved in the technical support of a PACS will require in-depth training. Those staff members who will be using the system intensively and using more of the facilities available (radiologists and radiographers) will require a greater degree of training than staff on wards. Note that there should be the facility to provide training for new staff as they join. Junior medical staff change every six months and it may be possible to include a short element of PACS training as part of any induction course.

16. Implementation

Staged implementation or 'big bang'.

Before implementation of a PACS can begin, the business plan must be in place, together with any required funding (see chapter 14, "Procurement Process"). Once this is finalised, the project should be publicised to the staff affected; the advantages for all staff of a PACS can be publicised in order to gain their support. It is important to give a realistic appraisal of the benefits and costs of PACS. The implementation of a PACS can be approached in a number of ways, depending on local conditions, working practices and available funding.

PACS implementation at an existing hospital has different problems to those for a new build. An existing site will have a film archive, and a plan should be developed for the acquisition of historical images from the film archive to the PACS via a digitiser. It is unlikely to be necessary to digitise the complete archive, but it will prove advantageous to have a proportion of the images acquired to PACS. For example, it may be decided that all images acquired in the previous year should be available for viewing. The facility should exist, however, for any of the existing films to be digitised on demand. The advantage of having a digital archive of the hospital's historical images is that the historical images are available for display alongside the newly-acquired digital images, avoiding the delays involved with digitising film on-demand, or the delays and problems of viewing film side-by-side with digital images.

A newly-built site will not have a film archive, and so it will be possible to have a "filmless" hospital from the time it opens. This will rely on the PACS being ready for full operation on time, however (at an existing site, should implementation of a PACS be delayed it will normally be possible to continue with conventional film-based operation). Installation of a computer network can be designed-in to a new build, whereas for an existing site the technical aspects of an imaging network's integration with existing networks must be considered, together with the impact of the work required for physical installation of the network cabling.

Consideration of the infrastructure required for the PACS must occur before the implementation starts. The various image acquisition devices have particular room requirements, as can the rooms where PACS reporting workstations are to be sited, and the more general areas where the clinical workstations will be situated. Printers and digitisers will all require space, together with power supplies and network connections. The computer room housing the main PACS servers and image data archive will have special requirements regarding ambient temperature and humidity and access control (the room should be easily accessible to the PACS technical support staff, e.g. the system manager, but should be locked so that unauthorised staff cannot gain access). Detailed information on PACS infrastructure requirements is available from the NHS Estates Agency (see chapter 19, "Web Links and Further Reading").

Implementation of a PACS can cover an entire imaging department, or can start with single modalities or clinical areas.

A phased introduction of PACS has advantages. The implementation process can pause at any point, to allow for re-evaluation and adjustment of future plans in the light of lessons learnt from earlier phases. This approach also allows for advances in technology to be incorporated into later stages of a PACS implementation project.

The transition from film-based to filmless operation can occur on a single date, or can be a phased transition, with both film production and acquisition of digital images occurring for a period of several months. The advantage of parallel film and digital operation is that staff can gain experience of the PACS in a less critical environment. There is also the opportunity to cure any problems with the PACS installation prior to filmless operation.

Staff training is crucial to the success of a PACS project. All staff requiring training should have an appropriate amount of time allocated as part of their normal working hours for this training. The training should not occur too far ahead of use of the new system, since this can lead to staff forgetting what they have learnt, and in some cases a delay can lead to disillusionment concerning the new system. Staff can be trained in a number of ways: direct teaching from a full-time trainer; teaching from a fellow staff member who will have been assigned to some training duties; or training by example during the working day.

17. Operational issues

System Management, fault reporting, consumables, system failures.

A PACS is a complex system; its routine management and use demands careful planning, and can require significant on-going investment in terms of staff support. Consideration must also be given to contingency plans in the event of failure of a component of the PACS or of the complete system.

■ Routine issues

There are a number of routine tasks that must be performed for the smooth operation of a PACS. These include creation of user accounts to enable staff members to access the system at a level appropriate for their requirements, management of system backups and the safe storage of backup media, training of new staff members to ensure that they can use the PACS to the level of skill required to perform their job effectively and efficiently, and quality assurance and quality control tasks, to ensure that the system is performing to the required standard. It will also be necessary to manage and monitor the interfaces between the PACS and external systems.

A PACS may require a supply of several items, including storage media for the archiving of image data, and processing chemicals should those be required for any film processors that are included as part of the PACS. Responsibility for the management of the supply of these items should be established.

Management of the PACS image data archive can be a complex issue, particularly when ensuring that all images that are likely to be required for viewing are available in on-line storage. The information needed for the management of the data archive can come from a number of sources, including a Hospital Information System and a Radiology Information System (see chapter 11, “Interfaces To External Systems”).

Another area that requires planning is the management of media in archive jukeboxes. If the intention is to remove media from a jukebox as the jukebox fills, to allow room for blank media to be loaded, then the manner in which the media are used becomes important. It is recommended that disks or tapes are loaded into the jukebox in small batches, adding a new batch as each preceding batch of media becomes full. In this way, any individual disk or tape will contain data acquired during a well-defined period, and so it will become easier to decide which can be removed.

Users may occasionally experience problems when using the PACS, sometimes due to a system problem and sometimes due to a lack of experience of the system on the part of the user. It is useful for a site to have a central point of contact for operational problems. This should be a person with in-depth knowledge of the PACS who will be able to identify whether a problem is a user problem or a system problem. A single point of contact is valuable since that person can rapidly become aware of system wide problems should a number of similar problems be reported simultaneously from around the site.

Problems

It is inevitable that problems will occur with a complex system such as a PACS. Contingency planning for such problems can minimise their impact and reduce the time required to recover from any problems. Failures can occur with individual components or with the system itself. Some possible areas of failure are listed below; this list should not be seen as exhaustive, as each PACS installation will have its own areas of concern.

Failure of an acquisition modality

Hospitals must consider what action to take in the event of an acquisition modality becoming unavailable, whether this is due to breakdown or to planned servicing. However, such actions are site-specific and are beyond the scope of this document.

Inability of an acquisition modality to send image data to PACS

It is possible that a modality can acquire an image, but then be unable to send the data to the PACS. This can be due to a fault at the modality, a fault with the interface between the modality and the PACS, a problem with the PACS network, a failure of the PACS storage, or a general failure of the PACS. Whatever the cause, it is important that the acquisition modality can store the image locally, at the point of acquisition. There are two reasons for this. The first is that it may be necessary to display the image soon after acquisition, and the second is so that once the fault is repaired, the image can be sent to PACS.

The local image store is ideally digital, either on fixed hard disk or on removable media such as optical disk or tape. If it is film, the images can be viewed on a conventional light box and acquired to PACS via a digitiser; however, this is not a recommended solution.

Inability of RIS or HIS to send information to PACS

There may be problems with external information systems (e.g. HIS or RIS) or with their interfaces to PACS. This can result in the information that PACS requires for its routine operation failing to arrive on the system.

Problems with RIS

If the PACS is configured so that it receives examination request information from a RIS, failure to receive this information will mean that an appropriate entry is not made in the PACS database, and the examination information will not appear on the acquisition worklist at the modality. In this case it will be necessary to enter manually the minimal amount of information required to uniquely identify the examination into the PACS. Once the external system or its interface is repaired, the manually created entry must be reconciled with the information delivered by the RIS.

Problems with HIS

If the PACS is configured so that it receives information from a HIS pertaining to patient location and clinic appointments, then a failure to receive this information will mean that the ability of the PACS to manage the on-line image storage will be compromised. For example, if the HIS sends the PACS information on impending clinic appointments, the PACS will be unable to ensure that images for those patients booked for clinic visits have been retrieved from archive storage to on-line storage. This can lead to delays whilst the required images are fetched in response to a user request.

Failure of a workstation or display monitor

Failure of a PACS workstation or display monitor is unlikely to cause more than minor inconvenience, since it is likely that equivalent workstations are located nearby.

Failure of PACS server and/or PACS network

A general system-wide failure will mean that newly-acquired images can not be sent to PACS. This is a global instance of the situation discussed above for the inability of an acquisition modality to send images to the PACS. It will also mean that images can not be sent to from the PACS for display at workstations.

18. Trends and future developments

ASP, immediate data availability, pay-per-image. Funding.

Currently in the UK, PACS data (both image data and patient data) is stored on-site. In recent years in the USA there has been a trend for healthcare institutions with a PACS to store their image data off-site (sometimes hundreds of kilometres away) at third party data repositories. The company providing this off-site storage is commonly referred to as the "Application Service Provider", or simply the "ASP." The advantage of this model is that the hardware and expertise required to manage the storage hardware are provided by the ASP, potentially saving the healthcare institution a large initial investment. Various funding models are available when using an ASP for storage of image data, including "pay per image" where a small charge is made for each image stored. As the network infrastructure develops, and network capacity becomes sufficient, this may become a viable model for UK healthcare institutions.

The ability to make all data available on-line, and available for immediate display, in a PACS has become feasible in the last two years. In this model, all of a site's images are available for immediate display from on-line storage. This has become possible due to the dramatic fall in cost and the increase in storage capacity of disk-based storage over the last few years.

Areas outside radiology are likely to use more of the services provided by a PACS, for example in radiotherapy treatment planning, where it is possible for the outline of a tumour drawn on a PACS system to be loaded into a radiotherapy treatment planning computer.

Home reporting may become more common, as high speed network connections and high quality display monitors become cheaper.

Integration of all forms of electronic records (PACS, RIS, HIS, voice dictation, EPR) is likely to increase.

Mergers between hospitals to form multi-site trusts are likely to increase the demand for PACS. In many cases, the merger may lead to particular clinical services being concentrated at one location, thus leading to the need to move images around from one site to another. The cost of transferring images stored on archive media such as film is significant, as is the loss rate.

Following reporting, automatic despatch of clinical reports to requesting GPs can be initiated by the PACS or the RIS. This could be via fax or email, for example.

19. Further reading

Links to many websites with useful PACS information, including PACS vendors and relevant professional bodies, can be found on the PACSnet web site at http://www.pacsnet.org.uk

Sources used in preparation of this report:

Books

J. Anthony Seibert, Larry J. Filipow, Katherine P. Andriole (eds), Practical Digital Imaging and PACS, American Association of Physicists in Medicine, Medical Physics Monograph No. 25 (Medical Physics Publishing, Madison WI, 1999), ISBN 0-944838-20-0

Keith J. Dreyer, Amit Mehta, James H. Thrall (eds): PACS – A Guide To The Digital Revolution, (New York: Springer-Verlag, 2002) ISBN 0-387-95291-8

HK Huang (ed): PACS: Basic Prinicples and Applications (Wiley-Liss, New York, 1999), ISBN 0-471-25393-6

Papers

IHE Technical Framework , Revision 5.0, October 2001.
http://www.rsna.org/IHE/tf/ihe_tf_index.shtml

R. Anthony Reynolds, Nicola H Strickland, The Hammersmith Project: A Large Scale PACS for a General Hospital, EuroPACS Newsletter, Volume 10 No. 1, December 1996.

R. Anthony Reynolds, Picture Archiving and Communication Systems, in "Diagnostic Radiology" ed. R.G. Grainger and D.V. Allison, Churchill Livingstone Publications, 1997

Horii, S.C. , Prior, F.W. , Bidgood, W.D.J. , Parisot, C. , Claeys, G. (1996 September 09. 21:32) . "DICOM: An Introduction to the Standard", RSNA 1993, available at
http://www.dicomanalyser.co.uk/html/introduction.htm

N H Strickland, Some cost-benefit considerations for PACS: A Radiological Perspective, British Journal of Radiology 1996, vol 69 pp1089-1098

ACR Standard For Digital Image Data Management, American College of Radiology, 1998-2001:
http://www.acr.org/cgi-bin/fr?tmpl:standards00,pdf:pdf/digital_image.pdf

Video Monitor Test Pattern Tutorials, Brigham and Women's Hospital, 1997:
http://brighamrad.harvard.edu/research/topics/vispercep/tutorial.html#bwh

Acknowledgements:

Thanks to Tahir Malik of the NHS Purchasing and Supplies Agency for chapter 14, "The Procurement Process."

The section on IHE in Chapter 9, "Standards and Formats", is based on the "IHE FAQ" written by Jonathan Turner (as part of the BIR IHE Implementation Group) for inclusion in the BIR "IHE Implementation Guide." This, in turn, was based on the "IHE FAQ" produced by the IHE group in the USA.

20. Glossary

PACSnet has published a glossary of PACS-related terms. The glossary includes terminology specific to PACS and medical imaging, and more general computing terms that are appropriate to discussions of PACS.

The PACS glossary is available from the MDA, report number MDA 02017

Contact details

PACSnet was established in the summer of 2001 and is part of the independent evaluation programme supported by the Medical Devices Agency (MDA) , an Executive Agency of the Department of Health. The team is based at St George's Hospital in London.

PACSnet provides independent technical evaluation of Picture Archiving and Communication Systems (PACS) and the components of such systems. Information provided by PACSnet is intended to assist purchasers of PACS with their purchasing decisions, implementation and efficient use of their PACS.

For further information on PACS and the work of PACSnet, please contact PACSnet at:

PACSnet
Bence-Jones Offices (Perimeter Road)
St George's Hospital
London SW17 0QT
www.pacsnet.org.uk

Mr. Jonathan E. Turner
Principal Evaluator
020 8725 2619
Jonathan.Turner@stgeorges.nhs.uk

Mr. Dewinder S. Bhachu
Head of PACSnet
020 8725 3315
Dewinder.Bhachu@stgeorges.nhs.uk

Mr. Alan R. McBride
Senior Evaluator
020 8725 3352
Alan.McBride@stgeorge.nhs.uk

For further information on the Device Evaluation Service, please contact:
Debbie Smith
Programme Manager
Room 1207
Medical Devices Agency
Elephant and Castle

London
SE1 6TQ
020 7972 8155
debbie.smith@doh.gsi.gov.uk

MEDICAL DEVICES AGENCY

MDA Evaluation Reports

MDA evaluation reports are published by the Medical Devices Agency, an Executive Agency of the Department of Health. They are available free of charge to the UK National Health Service (NHS), and are for sale to commercial organisations and other interested parties. A free catalogue of available reports can be obtained from the Orders Department, or downloaded from the MDA web site:

http://www.medical-devices.gov.uk

Ordering

Send your order to the address given below, stating the number, title and quantity of each report required. Your reports will be despatched by second class post the following working day. If you are not a representative of the NHS, you will be invoiced separately. Non-NHS customers are reminded that it is not possible to offer refunds for reports ordered in error.

Orders Department
Room 1207
Medical Devices Agency
Hannibal House
Elephant and Castle
London
SE1 6TQ

Tel: 020-7972 8181
Fax: 020-7972 8105
E-mail: dep@medical-devices.gov.uk

Enquiries

General publication enquiries should be directed to the Orders Department:

Tel: 020-7972 8181
Fax: 020-7972 8105
E-mail: dep@medical-devices.gov.uk

ISBN 1 84182 505 0

Smart number 36 26552 1