MEDINFO 98
B. Cesnik et al. (Eds)
Amsterdam: IOS Press

A hospital-wide distributed PACS based on Intranet

David Bandon, Yves Ligier, Gerhard Trayser, Christian Girard, Marianne Logean, Osman Ratib

Digital Imaging Unit, University Hospital of Geneva, Geneva, Switzerland

Abstract

A hospital-wide Picture Archiving and Communication System (PACS) is currently under development at the University Hospital of Geneva. After a first implementation including two one-terabyte optical libraries, the system is expanded to integrate all the imaging modalities of the hospital. The new storage requirement is 10 terabytes to cover three year archive.

A large distributed image archive has been designed including new archive servers for long-term storage and display servers for medium-term storage. The acquisition, archive and distribution cycles are performed using separated networks combining Fast Ethernet and Ethernet. Image files are distributed to the wide-hospital using a prefetching strategy or an Intranet server, RADIOLAB. The first mode takes advantage of the fully integrated hospital information system DIOGENE 2 to allow the automatic retrieval of studies in advance. The second mode provides a convivial study selection from any conventional WWW (World Wide Web) browser. Image files are then transmitted to the user's display station using HTTP (HyperText Transfer Protocol) and handled by OSIRIS software, which acts as a helper or viewer.

Such a system is expected to meet the time requirement, which is less than three seconds per image.

Keywords

PACS; Image management; Design; Intranet

Introduction

Since 1990, a PACS (Picture Archiving and Communication System) is under development at the University Hospital of Geneva. Its main characteristics are distributed storage architecture and a full integration to a large scale hospital information system, DIOGENE 2 [1].

A first system has proved the PACS feasibility in the hospital [2]. So far two one-terabyte optical jukeboxes are used to store images acquired from five imaging modalities. Image files can be visualized with an adapted image viewer, the OSIRIS software [3].

The project is now reaching a turning point in its evolution. The overall image modalities are planned to become digital in a four year period. In that context, the Geneva PACS is intended to provide a global storage solution and a wide image distribution to clinical wards. A step by step approach has been chosen to replace the radiology film library: the immediate step is to cancel the double hard copies performed for the CT and MRI modalities at the hospital, the ultimate step being a complete filmless hospital.

To achieve such goals, an improved design is needed. Therefore, we have chosen to expand the distributed storage strategy and allow image distribution through an Intranet server, RADIOLAB [4].

Design goals

The goals of the design are to provide operational image archive and retrieval services for a whole hospital with the following requirements:

- on-line storage for approximately three years of image data;
- sufficient speed for satisfactory image retrieval: the time for image access has to be less than 3 seconds for single image (such as chest image) which is a known limit [5] for PACS acceptability by clinical users;
- an evolution capability.

Figure 1 describes the acquisition modalities. All the modalities in the left panel are interfaced to the current PACS archive. The modalities on the right will be connected during the next four years. By this time the entire image acquisition will be digital. All these modalities are located in five different sites within two different buildings: the tomography section, the emergency room, the cardiology unit, the nuclear medicine unit and the conventional radiology unit.

The estimated storage requirements are 3 terabytes per year to be archived. Since the image files will be archived using a lossless compression providing an approximate compression ratio of 2, it is expected to manage approximately 6 terabytes on line.

Acquisition modalities		
Current modalities	Volume / year (GBytes)	Modalities to be interfaced
2 CT units	230	1 CT unit
1 MRI unit	60	1 MRI unit
2 US units	20	3 DSA units
1 CR unit	110	2 Nuclear medicine
		2 Fluoro rooms
		13 Xrays (film digitizer)

Figure 1 - Imaging modalities of the hospital

For financial reasons, it is not realistic to maintain six terabytes on the fastest storage devices. Therefore, a multi-level storage strategy has been defined. The latest is based on the recent study carried out in the Hospital of Geneva. Observations which have been confirmed by other hospitals [6] indicate that approximately 80% of all retrievals occurs within the first year from the date of the initial study. Therefore, we have decided to maintain a one year acquisition on magnetic storage to insure the shortest time response. In the same time, a three year acquisition will be kept on-line on slower optical disks. A last level is a manual archive for images older than three years.

First PACS system

Figure 2 is a diagram of the current operational PACS system at the Hospital of Geneva. The different steps such as acquisition, communication and visualization are performed by dedicated modules. After their acquisition, the images are converted to the DICOM/Papyrus format. This format which is fully compliant to the DICOM standard allows the storage of several images in a single file [7]. The newly acquired images are then transferred to a distributed archive server through a local Ethernet network using FTP (File Transfer Protocol). When this archive process is achieved, the image files can be retrieved from any display stations. The image file transfer is performed using a separated segment from the one of acquisition. This sharing of data traffic into two networks insures a better performance. Such an improvement has been proved with simulation studies [8]: it leads to a better communication between the PACS modules independently of the data acquisition and distribution workload.

Archive functionality

Two imaging sections are interfaced to the PACS storage: the emergency unit and the tomography/ultrasound section. Each of these sections has one or several imaging sources and has its own archive server. This strategy appeared to be the best solution to the geographical dispersion of the modalities within the hospital.

Each archive server is composed of a UNIX server station with a hierarchical storage system: a high-speed short term storage on a 2 gigabytes magnetic disk and a slow, low cost media in an optical disk juke box. The goal of such a hierarchy is to let the most recently acquired and most frequently retrieved studies on the magnetic disk.

A first 12 inch WORM jukebox, purchased in 1992, supports 800 gigabytes while the second, installed in 1994, handles 1.2 terabytes. The file management is operated by a commercial HSM (Hierarchical Storage Management) software producing a virtual UNIX file system, which spans the entire archive capacity.

Software for image retrieval

Two software packages manage one after the other the image retrieval [9]: ISIS and OSIRIS. The first one acts as a database browser. It provides a flexible tool for physicians to select images according to different criteria (patient name, birthday and time interval). When the selection is done, the images are downloaded to the user's display station and visualized with the OSIRIS software.

The PACS database itself was developed using the INGRES relational database management system (Relational Technology, Inc., Los Angeles, CA). It handles a minimal subset of data related to the images and radiological examinations. Moreover, it stores the reference of the optical disk on which the image files are stored.

Results

The current system has been tested in different clinical environments. One jukebox has recently been filled that results in the need for manual shelf management of the exported platters. No archived data has been lost and the image files can be easily retrieved. However, this system is not able to provide a satisfactory time response since a retrieval frequently exceeds one minute for an image set. Two reasons can explain this delay. First, since the magnetic disks have a restricted volume, image files may be frequently retrieved from the optical disks. This procedure results in poor time performance and requires regular swapping of the disks in the jukebox. Second, the network may saturate when the data workload increases.

New PACS system

For its new development, Geneva PACS has to provide an acceptable retrieval time (< 3 sec. per image) and has to manage a storage capacity of 10 terabytes. In order to satisfy these requirements, a large distributed design has been settled. Moreover, a prefetch strategy and an Intranet distribution allow an efficient image wide-distribution within the hospital.

Design

Figure 3 shows the PACS architecture currently developed. New archive servers are deployed in the different sections which have not already been equipped. In order to allow fast retrievals, a display server is put in each of these sections and in some clinical units. Such a server supports medium-term storage and enables a fast access for a cluster of display stations. A last server, the RADIOLAB server, is used for the image distribution through the WWW.

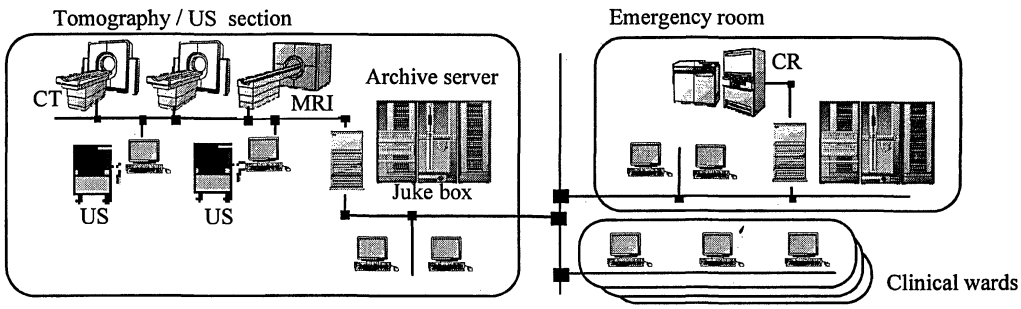

Figure 2 - Current PACS design

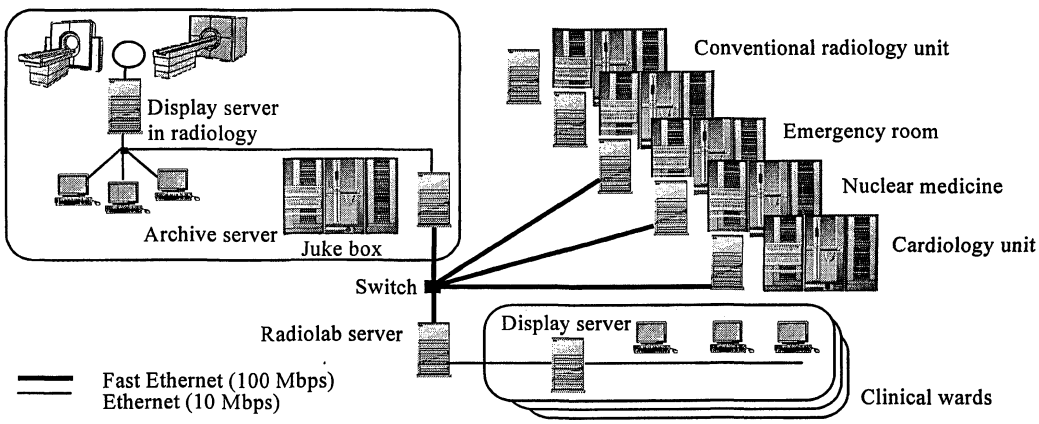

Figure 3 - New design being implemented

The last important evolution is the network improvement. Different high-speed networks are used for the image acquisition, archive and distribution.

Display servers

Two kinds of display servers are part of the new design.

The first one insures the best possible performance for primary diagnosis in Radiology (see fig. 3). A 20 Gigabyte RAID (Redundant Array of Inexpensive Disks) enables high-performance file reading. Moreover, a direct high-speed network connection allows a fast transmission to a cluster of display stations.

The image files are maintained on this server until their interpretation. After two weeks, if the diagnosis has not been performed, the image files are transferred to its archive server and deleted from the display server.

The second kind of display server is dedicated to image review in the clinical units. It includes a disk farm configuration with a capacity adapted to the needs of the clinical unit. Since physicians tend to move frequently from one location to the other and do not have a fixed workspace, they may retrieve images from any display stations. For that reason, the display server is expected to provide a fast access from any station located in the ward.

Archive servers

Four new archive servers are added in the design. They handle a similar hierarchical storage management as the two existing archive servers.

Network design

The new design takes advantage of a network upgrade currently underway within the hospital. This network is progressively moved from a bus to a tree topology (see figure 4). Therefore the new architecture is settled as a hierarchy combining Fast Ethernet in the upper levels and Ethernet in the lower levels. The Fast Ethernet protocol was selected because of its high bandwidth (100 Mbps) and its ability to deliver adequate performance. To avoid the saturation of the network, a switched design is incorporated.

Concerning the PACS, the previous hierarchical architecture is expanded with new segments entirely devoted to the image transmission. This modification guarantees a full bandwidth for the image transfer. Thus a new segment connects each display

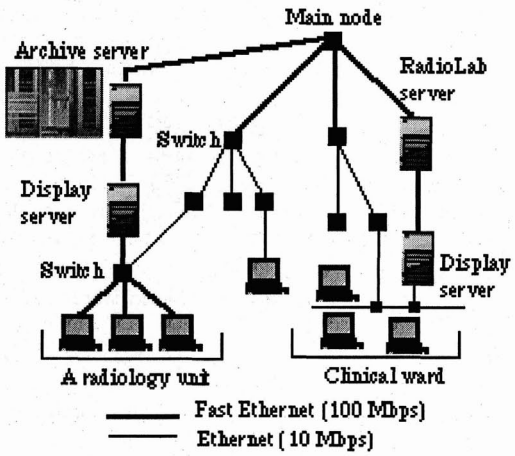

Figure 4 - Network configuration in the hospital

server of the radiology to its dedicated archive server. Moreover, archive and RADIOLAB servers are connected directly to the main node of the hospital (top node in the hierarchy). This structure allows for each archival server to be connected at the same time at the lower node related to its radiology area and at the top node. The first connection insures a direct image archive and retrieval for the radiology unit while the second one enables high speed image distribution to the clinical wards. Indeed, the image files go directly down in the network hierarchy from the top node.

Image retrieval

With the rapid growth of the PACS network and the increasing number of archive and display nodes, special software packages are needed to efficiently manage the transfer of image files between different servers. In fact image files must be sent from the archive node to the display servers where they are needed. Such a task is achieved by a traffic manager according to two different schemes:

- Prefetching allows the automatic retrieval of previous studies of a patient in advance. They are downloaded to a display server before the session of consultation occurs.
- A direct image selection provides image retrieval on demand. It is done using the RADIOLAB server which provides an efficient Intranet access to images through World Wide Web.

Prefetching

Such a strategy is based on observations of the medical activity. Radiologists and physicians usually compare newly acquired images with images from previous studies of the same patient. This comparison aims at evaluating evolution of the patient's medical state. Since this practice is known, prefetching consists in retrieving in advance the previous studies dealing with the patient. If these studies are too numerous, some rules are applied to select only the relevant studies in the patient's history.

The prefetch process is launched upon clinical events, such as scheduling, admission, and radiology registration. These various events are detected with the full integration between the different databases of the DIOGENE 2 system: IMPACT which manages the patients' admissions in the hospital and UNIMAGE which manages the examination schedule in the radiology. The link between these different databases is possible because a unique identification of all patients is applied.

In order to minimize the traffic at peak hours, older images are sent during night time when the traffic is at its lowest. However, if the patient is admitted in emergency, image files are immediately downloaded.

RADIOLAB

The RADIOLAB acts as database query and retrieval software. An interface coded in HTML is proposed from any conventional WWW browser on any platform (Mac, PC, UNIX.). The user can search studies in the patient records using different criteria, select some studies, preview some chosen iconized images and finally ask for their retrieval (see figure 5). This last step triggers a transfer by HTTP (HyperText Transfer Protocol) of images from a server to the local station. Image files are then opened, visualized and processed with the OSIRIS software,

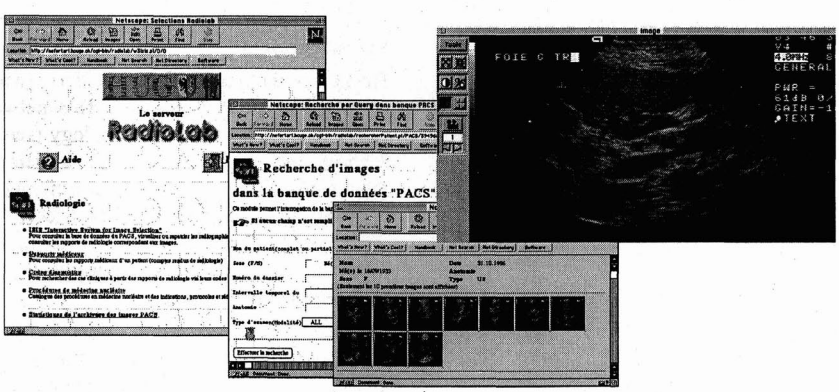

Figure 5 - User interface for image selection within RADIOLAB

which acts as a helper or viewer.

An optimization strategy is applied for the image retrieval. At a time, several copies of the same file can be stored on different servers. Therefore, the path allowing the highest-performance is used. Such a task is achieved by the traffic manager which stores in a dedicated database an updated list of image locations in the site.

Discussion

In this paper we propose a distributed architecture for a hospital wide PACS. The strength of our design is a centralized file management while the storage units are distributed. Coupled with an innovative technology and a high-speed network, this system is expected to provide a good time response (about three seconds per image).

RADIOLAB

The use of multiple display servers effectively improves the local image access and discharges the archive and network workloads. Besides, the data management is mainly performed on one server, the RADIOLAB server, where the database queries and the prefetch actions are concentrated. All the archive servers are managed on a same base using a hierarchical storage management. Therefore, a limited implementation effort has to be made and the usual difficulty of the distributed architecture is avoided: i.e. the very complex software tools for image management.

Another feature of our solution, which has to be considered in the evaluation of a solution, is the financial aspect. Since the storage costs decrease each year, such an approach is more realistic.

Moreover, a last interesting feature of the project is the full integration of the PACS in the large scale hospital information system, DIOGENE 2. A coherent link between all the data that constitutes the medical record (i.e., collections of notes, reports and observations) is offered. Therefore the hospital information system has now reached the point where the concept of computerized medical record becomes a reality. Software is currently implemented to provide a convenient and integrated access to all the multimedia data of the medical record. Images, which constitute an important part of a patient medical record, will be easily accessed through the PACS database with a unique identification for all the patients.

References

[1] J.R. Scherrer, R. Baud, D. de Roulet. Moving towards the future design of HIS: A view from the seventies to the end of the nineties, the DIOGENE paradigm. in: H.U. Prokosch, J. Dudeck, eds. *Hospital Information Systems: Design and Development Characteristics.* Amsterdam: Elsevier, 1995; pp. 347-375.

[2] O.Ratib, H.Lemke, G.Trayser, J.F.Vurlod, H.Do, J.R. Scherrer. Distributed image management and hierarchical storage in an integrated RIS and PACS. In: K.C. Lun, P.Degoulet, T.E. Piemme, O. Rienhogg, eds. *Proc. MEDINFO'92*, Amsterdam: Elsevier, 1992;pp 781-787.

[3] Y. Ligier, O. Ratib, M. Logean, C. Girard. OSIRIS: A medical image manipulation system. *MD Computing*, 1994: 11(4) pp 212-218.

[4] C. Mascarini, O. Ratib, G. Trayser, Y. Ligier, R.D. Appel. In-house access to PACS images and related data through World Wide Web. *European Journal of Radiology* 1996: 22 pp. 218-220.

[5] N.H. Strickland. Hammersmith PACS: Some lessons Learnt in Implementing a Filmless Hospital. *Proc. of the 14th International EuroPACS Meeting.*; Heraklion, (Greece) October 3-5 1996; pp. 12-17.

[6] M.M. Frost, J. Honeyman, E. Staab, PACS archive strategy utilizing a three-terabyte jukebox, in: *Medical Imaging 1996: PACS Design and Evaluation: Engineering and Clinical Issues.* R. G. Jost, D. J.Dwyer III, Eds. Proc. SPIE, 2711, pp 22-28 (1996).

[7] O.Ratib, H Hoen, C.Girard, Papyrus 3: DICOM-compatible file format. *Med. Inform.* 1994: 19(2) pp. 171-178.

[8] C.Debas, R.Kanz, J.F.Lerallut, J.Henry, D.Bandon, G.Evers, J.C.Kurdziel. Introduction of modelling & simulation in the lux-imacs project. *Proc. of the 12th International EuroPACS Meeting*, Geneva, September 22-24, 1994.

[9] Y.Ligier, O. Ratib, C. Girard, M. Logean, G. Trayser. Distributed file management for remote clinical image viewing stations, in: *Medical Imaging 1996: PACS Design and Evaluation: Engineering and Clinical Issues.* R.G. Jost, D.J.Dwyer III, Eds. Proc. SPIE, 2711, pp 475-482 (1996).

Address for correspondence

David Bandon
E-mail: David.Bandon@dim.hcuge.ch
Digital Imaging Unit
University Hospital of Geneva,
24 rue Micheli-du-Crest
1211 Geneva
Switzerland.