Transparent Image Access in a Distributed Picture Archiving and Communications System: The Master Database Broker

R.D. Cox, C.J. Henri, and R.K. Rubin

A distributed design is the most cost-effective system for small- to medium-scale picture archiving and communications systems (PACS) implementations. However, the design presents an interesting challenge to developers and implementers: to make stored image data, distributed throughout the PACS network, appear to be centralized with a single access point for users. A key component for the distributed system is a central or master database, containing all the studies that have been scanned into the PACS. Each study includes a list of one or more locations for that particular dataset so that applications can easily find it. Non-Digital Imaging and Communications in Medicine (DICOM) clients, such as our worldwide web (WWW)-based PACS browser, query the master database directly to find the images, then jump to the most appropriate location via a distributed web-based viewing system. The Master Database Broker provides DICOM clients with the same functionality by translating DICOM queries to master database searches and distributing retrieval requests transparently to the appropriate source. The Broker also acts as a storage service class provider, allowing users to store selected image subsets and reformatted images with the original study, without having to know on which server the original data are stored.

Copyright © 1999 by W.B. Saunders Company

THE ULTIMATE GOAL in the design of picture archiving and communications systems (PACS) and image management systems is to bring image data to the user in a timely and efficient manner. This presents a challenge to designers of a distributed system, where images could be located on any of several different PACS servers having different characteristics for their storage media. Users of the PACS should not be required to know where the data are located when they search for and retrieve images.

PACS INFRASTRUCTURE

The PACS at the McGill University Health Centre (MUHC) was designed and implemented in-house¹ using freely available software and development tools. We chose to conform to a distributed model, with the one exception being a centralized patient listing or Master Database. This design has proven cost-effective while still allowing filmless radiology.

For the purpose of the PACS design, we view

image storage as hierarchical, with three broad categories or types: (1) on-line—images are on a medium available for immediate access such as a disk subsystem; (2) near-line—images are located on a medium that can be accessed in an automated fashion, but are not immediately available to the requesting application. Examples would be optical jukeboxes and magnetic tape libraries; and (3) off-line—image access requires manual intervention such as the loading of an optical disk, tape, or CD.

Each of our PACS storage components is defined as one of the above types. They include the following: (1) Modality Server—provides on-line storage of image data sent from scanners; (2) DLT Archival/Retrieval Server—provides near-line storage on digital linear tape (DLT). Images are transferred from the Modality Servers as they are received; (3) CD Archival Server—provides a facility for off-line storage by “burning” image data on recordable CD; and (4) CD Retrieval Server—used to retrieve off-line data from CD. Both the DLT and CD Retrieval Servers have a retrieval cache that stores recently retrieved image data. These images are considered on-line as long as they are in the cache. All storage components are Digital Imaging and Communications in Medicine (DICOM)-conformant, providing query/retrieve and storage services.

In addition, the PACS consists of devices that do not provide image storage. Nonstorage components include: the Master Database Server, which contains a record of all scanned digital data, and the Web Server, which provides web-based access to images and radiology reports.

Our current PACS installation consists of six Modality Servers, including two computed tomography (CT), one magnetic resonance (MR), one

From the Department of Diagnostic Radiology, McGill University Health Centre, Montreal, Quebec, Canada.

Address reprint requests to Robert D. Cox, BEng, Department of Diagnostic Radiology, Montreal General Hospital, 1650 Cedar Ave, Montreal, Quebec, Canada H3G 1A4.

*Copyright © 1999 by W.B. Saunders Company
0897-1889/99/1202-1052\$10.00/0*

ultrasound (US), and two nuclear medicine (NM) servers (one for our NM department and one for the NM department at an affiliated hospital). We have one DLT Archival/Retrieval server with an ADIC Scalar 458 tape library (Advanced Digital Information Corp, Redmond, WA), one CD Archival Server, one CD Retrieval Server, two Master Database Servers (the second acts as a redundant system), and one Web Server.

Images remain on-line on the Modality Server that received them until they are deleted to make room for newer data. The amount of on-line storage is a function of the size of the disk subsystem (redundant array of inexpensive disks [RAID]) installed and the volume of data scanned, but is typically 6 months. Our current near-line storage device is estimated to hold about 3 years of data with the remainder available as off-line storage. Off-line storage capacity is essentially infinite.

The design approach we have used is to distribute the images while keeping a central record of their current location. The Master Database contains tables for patient, study, and series levels arranged in a hierarchical fashion with unique keys to access records in each table. The structures of these tables provide the information required for DICOM query/retrieve. A fourth table is included (using the same unique key as the series level) that lists one or more locations for each series within the PACS, where "location" could be any of the on-line, near-line, or off-line storage devices. Other locations, such as DICOM workstations and the Web Server, are considered PACS clients and are not listed.

Ideally, clients query the Master Database, which lists all patients and studies regardless of their location, and the server takes care of directing the client to the correct location when data are requested. Our PACS browser client, which was developed in-house,² was designed to provide this functionality by using a distributed web-based viewing system. However, to provide the same level of service to other clients, such as commercially developed DICOM workstations, we needed to develop a standards-based solution.

MASTER DATABASE BROKER

The Master Database Broker runs on the Master Database Server and was developed using the Mallinckrodt Institute of Radiology Central Test

Node (CTN) code library,³ which provides a framework for creating DICOM-conformant applications.

The design for the query/retrieve module of the Broker includes both Patient Root and Study Root information models. That is, clients can request patient and study information separately or as a single list. This feature was necessary as we have workstations that support each of these DICOM query/retrieve models. In response to a C-FIND request from a DICOM client, the Broker queries the Master Database using the search criteria in the request and returns the resulting patient/study/series lists to the client.

When a case is retrieved from the Broker, the request is processed and an internal list of series required for the transfer is produced. Then, for each series in this list, the Broker determines the optimal location from which to request that series, using the same algorithm used by the PACSbrowser client. Specifically, the Broker examines each location listed for the series, pulling the Application Entity Title (AE Title). The Broker then consults a master AE Title database (called GlobalCTNCtrl), which designates whether the server associated with that AE Title is an on-line, near-line, or off-line storage device. If the AE Title matches the criteria for on-line storage, a C-MOVE is initiated to that AE with the client that requested the retrieval as destination. If no on-line match is found, the Broker checks GlobalCTNCtrl for a near-line location. A match initiates a C-MOVE to the near-line storage device. The series is retrieved from the near-line medium, for example DLT, and sent to the requester. Failure to find a near-line location triggers a request to the off-line storage device, which requires manual intervention to fulfill the request, for example, the insertion of a particular CD into the CD-ROM drive of the CD Retrieval Server.

The Broker also acts as a storage service class provider for secondary capture and reformatted images created on workstations. Each study in the PACS is coded with an "Owner" AE Title, typically the Modality Server that first received the image data. This key is used by the Broker to transfer any images it receives to the same location as the rest of the study. In this way, users do not need to know to which server to send a series. They simply send the images to the Broker, which

forwards them to the appropriate storage device using a DICOM “push.”

CONCLUSION

The Master Database Broker has been operational since October 1998, giving DICOM workstation clients the same PACS-wide “view” of patient image data previously only available to locally developed applications such as the web-based PACSbrowser.

With the addition of the Broker, we have been able to retain all the advantages of a distributed

PACS, while providing a centralized system from the perspective of users. This has greatly enhanced the usability of our PACS implementation.

REFERENCES

1. Cox RD, Henri CJ, Rubin RK, et al: DICOM-compliant PACS with CD-based image archival. SPIE 3339:135-142, 1998
2. Henri CJ, Rubin RK, Cox RD, et al: Design and implementation of WWW-based Tools for image management in CT, MRI and ultrasound. Can J Med Radiat Technol 28:135-138, 1997
3. Moore SM, Hoffman SA, Beecher DE: DICOM shareware: A public implementation of the DICOM standard. SPIE 2165:772-781, 1994