Finding the Optimal Picture Archiving and Communication System (PACS) Architecture: A Comparison of Three PACS Designs

Wyatt M. Tellis and Katherine P. Andriole

Purpose: At present, there are two basic picture archiving and communication system (PACS) architectures: centralized with a central cache and controller, and distributed with a distributed cache and central controller. A third architecture proposed here is an autonomous one with a distributed cache and no controller. This report will investigate the performance (as measured by central processing unit [CPU] and network load, scalability, and examination retrieval and display latency) of these three types. **Methods:** The distributed PACS architecture will be simulated using an IMPAX R3.5 (AGFA, Ridgefield Park, NJ) PACS, while the centralized design will be simulated using an IMPAX R4 (AGFA) PACS. The autonomous system will be designed and implemented in-house. The autonomous system consists of two types of entities: basic components such as acquisition gateways, display stations, and long-term archives, and registry servers, which store global state information about the individual PACS components. The key feature of the autonomous system will be the replacement of the central PACS controller by the registry servers. In this scenario the registry servers monitor the interactions between the components, but do not directly govern them. Instead each component will contain the application logic it requires and will use the state information from the registry servers to take the appropriate action, such as routing images, prefetching studies, and expiring images from near line cache. In addition the routing of examinations will be optimized to reduce the duplication of image data. Display stations will be categorized by specialty (neuroradiology, pediatrics, chest, etc) and will retrieve studies for display on demand from intermediate servers dedicated to the corresponding specialty. Studies will be routed only to the intermediate servers and not to display stations. **Results:** By distributing the application logic, an autonomous PACS architecture can provide increased fault tolerance and therefore increased uptime. In addition, the lack of a central controller and the use of intermediate servers

improve the scalability of the system, as well as reduce CPU and network loads.

Copyright © 2001 by W.B. Saunders Company

WITH THE INCREASED use of picture archiving and communication systems (PACS) in radiology departments throughout the world, system performance has become an important issue in the transition to filmless radiology.^{1,2} This project will analyze two PACS architectures, a centralized cache system and distributed cache system, to identify the limiting factors in performance. The final goal is to construct a hybrid system that incorporates the benefits of each architecture while eliminating the bottlenecks.

The basic study management operations a PACS must perform are query, retrieve, and routing. The query operation is defined as the searching of a database of examinations using various criteria, such as patient name, medical record number or accession number. The actual database of examination information may be centralized, with a single independent database, or distributed, with multiple independent databases each responsible for its own set of examination information. A retrieve operation is defined as the "pulling" of images from one storage location, such as long-term, near-line, or local storage, to another. This operation may use the Digital Imaging and Communications in Medicine (DICOM) C-MOVE or C-GET operations, the file transfer protocol (FTP), or proprietary mechanisms. The routing operation is similar to retrieval except it involves "pushing" images from one location to another.³

This project will compare two PACS with different image cache architectures: a distributed image cache system versus a centralized cache system. In each architecture there is a centralized PACS controller that tracks the location and manages the transfer of images throughout the PACS. In addition, there are three basic types of devices: an archive server, an acquisition gateway, and a display station. In the distributed cache environment, the image cache is distributed across all three types of devices, which allows an image to be retrieved from any device that currently contains it. In such a system when new images arrive at the

From the Biological and Medical Informatics Graduate Program, and the Department of Radiology, University of California San Francisco, San Francisco, CA.

Address reprint requests to Wyatt M. Tellis, BA, Laboratory for Radiological Informatics, University of California San Francisco, 530 Parnassus Ave, Room CL-158, San Francisco, CA 94143-0628. E-mail: wyatt.tellis@radiology.ucsf.edu.

Copyright © 2001 by W.B. Saunders Company

0897-1889/01/1402-1018\$35.00/0

doi:10.1053/jdim.2001.23822

Table 1. Hardware Configuration of Distributed PACS

Type	Count	CPU	RAM	Storage	Network
Oracle server	1	6 × 366 MHz Ultra SPARC 2	2 GB	108 GB mirrored	155 Mbit ATM
Archive server	1	2 × 200 MHz Ultra SPARC 2	256 MB	500 GB mirrored	155 Mbit ATM
Gateway	3	2 × 200 MHz Ultra SPARC 2	128 MB	13 GB	155 Mbit ATM
Display station	1	1 × 143 MHz Ultra SPARC 2	128 MB	13 GB	155 Mbit ATM
	1	1 × 167 MHz Ultra SPARC 2	256 MB	13 GB	155 Mbit ATM
	1	1 × 200 MHz Ultra SPARC 2	128 MB	9 GB	155 Mbit ATM
	1	1 × 333 MHz Ultra SPARC 2	256 MB	18 GB	155 Mbit ATM
	1	2 × 200 MHz Ultra SPARC 2	128 MB	20 GB	155 Mbit ATM
	1	2 × 200 MHz Ultra SPARC 2	192 MB	13 GB	155 Mbit ATM
	1	2 × 200 MHz Ultra SPARC 2	256 MB	20 GB	155 Mbit ATM
	1	2 × 300 MHz Ultra SPARC 2	192 MB	13 GB	155 Mbit ATM
	1	2 × 300 MHz Ultra SPARC 2	256 MB	22 GB	155 Mbit ATM
	1	2 × 300 MHz Ultra SPARC 2	256 MB	30 GB	155 Mbit ATM
	2	1 × 167 MHz Ultra SPARC 2	128 MB	9 GB	155 Mbit ATM
	3	2 × 200 MHz Ultra SPARC 2	128 MB	13 GB	155 Mbit ATM
	4	1 × 200 MHz Ultra SPARC 2	128 MB	13 GB	155 Mbit ATM
	4	2 × 300 MHz Ultra SPARC 2	256 MB	13 GB	155 Mbit ATM

acquisition gateways they are automatically routed to the appropriate display stations and archive server. In addition, each display station has its own local cache, which is used to store a local copy of an image for instantaneous display. As the name suggests, the centralized architecture has a single, centrally located image cache. In this architecture newly acquired images are kept in the central cache and never directly routed to the display stations. Instead each display station pulls images from this central cache on demand whenever an image is displayed.

MATERIALS AND METHODS

The distributed cache PACS is an IMPAX R3.5 PACS. It consists of one central Oracle server running the Oracle 7 DBMS (Oracle Inc, Redwood Shores, CA), one archive server with a 1-TB magneto-optical disk (MOD) jukebox, three acquisition gateways, and 23 display stations. All devices within the PACS are Sun SPARC computers running the Solaris 2.5 operating system (Sun Microsystems, Mountain View, CA). The hardware configuration of each type of device is detailed in Table 1.

The centralized cache PACS is an IMPAX R4 PACS (AGFA, Ridgefield Park, NJ). Like the R3.5 architecture, the R4 architecture contains the four basic types of devices: an Oracle server running Oracle 8, which acts the PACS controller, an archive server with attached MOD jukebox, three acquisition gateways, and multiple display stations. Unlike R3.5, the R4 system is a mix of Sun SPARC computers running Solaris 2.7 and Dell (Round Rock, TX) PCs running the Windows NT operating system (Microsoft, Redmond, WA). The proposed configuration of the R4 PACS is detailed in Table 2. With the R4 PACS there will be an increase in the size of the cache, but a decrease in its distribution across the PACS. With this architecture only the archive server and three gateways will have local image storage.

While testing of the full centralized cache PACS has been deferred until delivery of the clinical system, several performance measurements were made of the distributed cache PACS. These include total image transfer load, acquisition transfer load, and acquisition routing times. The total image transfer load was calculated by tracking the number of images that arrived at a device within a given hour. The acquisition transfer load was calculated in a similar manner as the total image load, but only images from studies done on the current day were included in the summation. The acquisition routing

Table 2. Hardware Configuration of Centralized PACS

Type	Count	OS	CPU	RAM	Storage	Network
Oracle server	1	Solaris 2.7	6 × 366 MHz Ultra SPARC 2	2 GB	400 GB RAID 5	Gbit Ethernet
Archive server	1	Solaris 2.7	4 × 480 MHz Ultra SPARC 2	4 GB	720 GB RAID 5	Gbit Ethernet
Gateway	3	Solaris 2.7	4 × 480 MHz Ultra SPARC 2	4 GB	720 GB RAID 5	Gbit Ethernet
Display station	~25	Win NT 4.0	Pentium III	256 MB	10 GB	Gbit Ethernet

Fig 1. System-wide results (all transmission times in seconds).

time was calculated by measuring how long it took to route newly acquired images from a gateway to their destination devices. The statistics were captured over a 28-day period during the month of December.⁴⁻⁷

RESULTS

The results of our analysis of the distributed PACS can be found in Figs 1 through 3. Figure 1 contains the aggregate statistics for all devices (excluding the Oracle server and gateways) within the PACS. Figure 2 shows the results for the archive server, and Figure 3 contains the results for the acquisition gateways. It should be noted that Figure 3 does not contain data for acquisition transmission times since the gateways are the initial destination for images and therefore would have a transmission time of zero.

Two trends stand out in our analysis, both of which are especially evident on the archive server.

The first is the increase in total load as measured by megabytes per hour and images per hour during the early morning and midday hours. The increase during the early morning can be explained by an increase in "pre-fetch" activity. The early morning is the time during which most of the prior studies for examinations scheduled for that day are pre-fetched from the long-term archive and routed to the appropriate display station. The midday increase in total load can be explained by an increase in manual retrieval from the long-term archive. Such manual retrieval usually occurs when users wish to view old studies for research, teaching, and comparison.

The second trend was a spike in the transmission times of newly acquired images during the morning hours. This exact cause of this increase is still being investigated; however, it is believed the daily

Fig 2. Archive server results (all transmission times in seconds).

Fig 3. Gateway results.

backup of the Oracle database, which occurs at this time, may be contributing to the phenomenon. It is hypothesized that during the backup, the Oracle server, which oversees the transfer of images across the PACS, commits some of its computing resources to the backup operation, effectively reducing its overall capacity to handle database activity.

DISCUSSION

Each architecture has its own advantages and disadvantages. With a centralized PACS the process of routing images is eliminated since there is only one true cache location. This greatly simplifies the design of the PACS and reduces the load on the PACS controller since it no longer has to track the location of images across multiple caches. The elimination of routing, however, may result in larger spikes in network

traffic, since routing in combination with pre-fetching allows for the image transfer load to be distributed over a greater time period thus maintaining a more constant network load. In a distributed cache architecture an image can be pre-fetched and routed to the local cache on a display station during periods of low usage, such as the early morning. This eliminates the need for the real-time retrieval of image data from a central cache whenever an image is displayed. Since it appears there is considerable image transfer activity during the midday (Fig 1), a time which coincides with increased usage, having the additional network traffic associated with the real-time retrieval of images may result in poorer display station performance in a centralized PACS. In contrast an increase in midday activity has no effect in a distributed PACS since all

Fig 4. Cluster cache PACS architecture.

images are stored locally on a display station where the process of displaying an image only involves reading it from disk.⁶

In addition, even though both systems have a central PACS controller the centralized cache architecture introduces an additional central point of failure. Since all images are stored in a single cache location, the failure of this cache brings all operations within the PACS to a halt. This contrasts with the distributed system where the failure of one cache can be compensated for by the use of other caches that contain copies of images stored on the failed cache.⁶

Given the advantages and disadvantages of the two systems, an architecture that combines the two might prove to be an ideal solution. A system that uses a clustered cache approach with cacheless display stations serviced by individual image servers takes advantage of the best features of a distributed and centralized PACS (Fig 4). In a clustered architecture images are routed to the image servers, thereby distributing network traffic across a greater time period. Images, however, are not routed directly to the display

station, which reduces network traffic by reducing the number of copies of an image that must be sent. In addition, a display station only pulls the images that are to be displayed, further reducing network traffic. Since only a subset of all display stations retrieve their images from a particular server, the spike in network traffic associated with the real-time transmission of images will not be as great as when all display stations are retrieving from one location, as is the case in a centralized PACS. Each image server can be set up to primarily service the display stations within a particular section or floor. This means only images of interest to the section associated with the image server need be sent to the server. This policy is similar to the current routing rules used to route images in the distributed PACS. In addition, a second copy of an image can be routed to an additional, unrelated image server. This second copy can act as a backup for when a display station's primary image server fails. In such an event the display station can be instructed to retrieve the copy from the station's designated backup server.

REFERENCES

1. Arenson RL, Andriole KP, Avrin DE: Computers in imaging and health care: Now and in the future. *J Digit Imaging* 13:145-156, 2000
2. Bauman RA, Gell G: The reality of picture archiving and communication systems (PACS): A survey. *J Digit Imaging* 13:157-169, 2000
3. Huang HK: PACS: Basic Principles and Applications. New York, NY, Wiley-Liss, 1999
4. Richardson, NE, Thomas JA, Lyche DK, et al: The philosophy of benchmark testing a standards-based picture archiving and communications system. *J Digit Imaging* 12:87-93, 1999
5. Shelton PD, Lyche DK, Norton GS, et al: Benchmark testing the digital imaging network-picture archiving and communications system proposal of the Department of Defense. *J Digit Imaging* 12:94-88, 1999
6. Wong AWK, Lou SL, Huang HK: PACS subsystems performance: Centralized versus distributed. *SPIE* 1899:24-30, 1993
7. Wong AWK, Huang HK: Subsystem throughputs of a clinical picture archiving and communications system. *J Digit Imaging* 5:252-261, 1992