Enterprise-Class Digital Imaging and Communications in Medicine (DICOM) Image Infrastructure

Gary York, Joe Wortmann, and Razvan Atanasiu

Most current picture archiving and communication systems (PACS) are designed for a single department or a single modality. Few PACS installations have been deployed that support the needs of the hospital or the entire Integrated Delivery Network (IDN). The authors propose a new image management architecture that can support a large, distributed enterprise. Copyright © 2001 by W.B. Saunders Company

WHEN Digital Imaging and Communications in Medicine (DICOM) image management is provided to an entire Integrated Delivery Network (IDN) as a fundamental infrastructure service, a number of critical issues must be addressed, including:

- Scalability and performance
- Distributed capture, storage, and viewing
- High availability
- Ease of management, support, and upgrade
- Standard support, openness, and integration with other systems

DISTRIBUTED ARCHITECTURE

The new architecture is based on a high-performance archive manager that can scale to handle the needs of the enterprise, along with distributed DICOM caches that can handle local needs for storage and access. Figure 1 shows an implementation with a single central archive and multiple distributed caches.

The DICOM archive manager is designed as a high-performance, scalable archive for DICOM objects (images, structured reports, and presentation state). It serves as a central repository for DICOM objects so that the cost of management and support can be kept low. High availability is achieved through the use of redundancy and automatic failover. Performance is achieved through the efficient use of machine resources and support for a distributed archive. As the enterprise load increases, additional machines can be added to a high-availability cluster. By avoiding centralized resources that can lead to contention and performance bottlenecks, the archive manager can continue to have excellent performance as new workstations and modalities are brought on-line. The archive manager also supports efficient utilization

of machine resources through DICOM image streaming and a thread per association policy.

ARCHIVE MANAGER

Scalability and Performance

An enterprise class archive manager must have excellent scalability and performance. Some cardiology modalities can be substantially larger than typical radiology modalities. The archive manager must be designed to handle these large modalities without degrading the performance of other users of the archive. The authors have implemented a new streaming model for DICOM communications that allows very large modalities to be received by the archive and sent from any DICOM storage class user (SCU) or storage class provider (SCP) using a small, configurable, fixed amount of memory. This allows the archive manager and DICOM caches to send and receive extremely large studies with controlled machine resources.

The implementation also uses a very efficient *thread per association* model for DICOM communications. Each new DICOM association uses an operating system thread, instead of using a new process. Since operating system threads can be an order of magnitude more efficient than processes, this reduces the memory requirements and decreases the context switching time between associations. Combined with database and DICOM association connection pooling, performance and scalability are improved. We have demonstrated an archive manager that handles over 300 active, concurrent DICOM associations on a two-processor, departmental server.

The archive manager allows multiple DICOM application entities to be created on multiple servers in a server cluster. This allows for the simple addition of a new server to the cluster to handle increased enterprise load.

From Emageon, Inc, Birmingham, AL

Address reprint requests to Gary York, PhD, Emageon, Inc, 1200 Corporate Dr, Suite 400, Birmingham, AL 35242. E-mail: gyork@emageon.com.

Copyright © 2001 by W.B. Saunders Company

0897-1889/01/1402-1016\$35.00/0

doi:10.1053/jdim.2001.23820

Fig 1. Central archive and distributed caches.

System Management

Enterprise class applications also must be easy to manage, and must be manageable while the system is running. The archive manager and the DICOM caches can be reconfigured while the system is running. A system administrator is also able to monitor activity while the system is running. A central, remote administrative graphical user interface (GUI) can be used to monitor the current activity, review the activity logs, and re-configure the properties of any archive manager or DICOM cache.

High Availability

High availability must be a core capability of an enterprise class image manager. Some degree of high availability can be achieved through off-the-shelf high-availability hardware and software. Clustering and redundancy can improve the likelihood of the system surviving single points of failure. Beyond these capabilities, the system should be designed to provide high availability. The archive manager can be combined with distributed DICOM caches to provide for image availability in the presence of component failures.

DICOM CACHE

The DICOM archive teams with any number of distributed DICOM caches to provide the image service to the enterprise. The caches can be physically dispersed throughout the enterprise to take advantage of local creation and access patterns. Each cache acts like the full, enterprise DICOM

archive; however, only a subset of the images will actually be stored in the cache at any time. By providing local access to the images that are most likely needed, the load on the central archive can be reduced, thereby providing better overall system performance. In addition, the distributed model improves availability since downtime of one component does not affect the availability of other components in the system. The DICOM caches provide a "proxy service" that allows multiple back end archives to be consolidated into a single view. This allows the enterprise to support multiple DICOM archives, while providing a consolidated view of all images in the system. The caches provide a "write through" capability for images sent from a modality. Such images are scheduled for routing to an archive using a rule-based, reliable delivery router. Image delivery can be scheduled to take advantage of low usage periods on the local and wide area networks (LAN and WAN).

Proxy

The proxy service provides access to the contents of the cache and to one or more underlying DICOM storage providers. The proxy supports DICOM standard C-Find, C-Move, C-Get, and C-Store capabilities. When the proxy is queried, the contents of the cache are combined with the contents of each of the underlying DICOM storage providers. The proxy communicates with the underlying storage devices using DICOM. The responses of each query are combined to eliminate duplicates, and the response to the query is completed. In order to improve query performance, the proxy orders the queries, and incrementally replies to the query as each underlying device responds. The underlying storage systems can contain duplicates, which are eliminated as the reply to the query is formed.

The C-Move operation can also be performed through the proxy. If the DICOM object is not in the cache, the proxy still transparently retrieves it. The proxy will attempt to complete the move using an ordering of the underlying storage providers. Failure will occur only if all of the underlying systems cannot complete the retrieval.

The DICOM cache architecture supports a distributed high-availability environment. First, if a cache is not available, the clients to the cache can make requests directly to the underlying archive manager(s). Second, the cache is designed to provide partial responses to requests if it cannot com-

municate with an underlying DICOM storage provider. At a minimum, a cache client can have access to all the images in the cache even in the network of archive manager is not available. Finally, if the underlying DICOM storage devices provide duplicate access to the same studies, even if one provider is unavailable, the proxy will skip the unavailable provider and move on to the duplicate provider.

Cache Management

The caches are actively loaded with images by a ruled-based prefetcher. Based on events such

as examination scheduling, modality worklists, and performed procedure steps, images can be prefetched from an archive to the cache. By using intelligent prefetching, we aim to allow the caches to handle more than 95% of the requested image load.

The content of the caches is managed through a *least recently used* (LRU) deletion strategy. When the cache is full, items that have not been recently used will be deleted. The deletion strategy is extremely flexible, allowing improved strategies to be used based on site-specific information.