

AD _____

Award Number: DAMD17-00-1-0585

TITLE: Reliability, Security, and Authenticity of Meta Medical
Image Archive for the Integrated Healthcare Enterprise

PRINCIPAL INVESTIGATOR: Robert G. Gould, Sc.D.

CONTRACTING ORGANIZATION: University of California, San Francisco
San Francisco, CA 94143-0962

REPORT DATE: October 2004

TYPE OF REPORT: Final Addendum

20060223 004

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Approved for Public Release;
Distribution Unlimited

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE October 2004	3. REPORT TYPE AND DATES COVERED Final Addendum1 (1 Oct 2003 - 30 Sep 2004)	
4. TITLE AND SUBTITLE Reliability, Security, and Authenticity of Meta Medical Image Archive for the Integrated Healthcare Enterprise			5. FUNDING NUMBERS DAMD17-00-1-0585	
6. AUTHOR(S) Robert G. Gould, Sc.D.				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California, San Francisco San Francisco, CA 94143-0962 E-Mail: Robert.gould@radiology.ucsf.edu			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Medical Research and Materiel Command Fort Detrick, Maryland 21702-5012			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) Not Provided				
14. SUBJECT TERMS Not Provided			15. NUMBER OF PAGES 23	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unlimited	

Table of Contents

Cover.....	1
SF 298.....	2
Table of Contents.....	3
Introduction.....	4
Body.....	4
Key Research Accomplishments.....	18
Reportable Outcomes.....	18
Conclusions.....	19
References.....	none
Appendices.....	20

Introduction

This project involved the design, development and testing of a prototype Meta-Medical Image Archive (MMIA) that is able to facilitate the electronic exchange of diagnostic imaging studies between institutions. The overwhelming majority of imaging studies are done with digital modalities and the images are distributed, viewed and stored electronically without the use of film. When a patient receives care at more than one facility, electronic transfer of image data between institutions would benefit patient care but such transfer must not interrupt local operations and must ensure the integrity of the information. An MMIA design was developed and implemented in a testbed environment that uses a clustered component configuration with full failover capabilities. The MMIA enables a user at an image display device of a Picture Archiving and Communications Systems (PACS) at one institution to query and identify studies stored either within the archive of the MMIA or in an archive of a PACS at a different institution that connects to the MMIA, and then retrieve and view those studies at the local workstation. The user directly interacts only with the MMIA and there is no interference with the local PACS at any institution. Data transfer between the MMIA and the local PACS can be encrypted to assure integrity and safety of the transfer of the information. This project implemented an MMIA in a controlled environment and evaluated its ability to query multiple PACS, to retrieve, forward and store studies, and to failover in the event of component failure. The impact of encryption on data transfer times was also investigated.

Body

1) Project planning (Task 1).

The MMIA design was developed to address scenarios where an individual receives healthcare at multiple institutions. It was assumed that all institutions have digital imaging capabilities and have implemented Picture Archiving and Communication Systems. No limitations were imposed on the distance between or geographic location of institutions but the assumption was made that all connect to a common network, presumably the internet, although the MMIA could be implemented over a private network as well. No requirements were placed on network bandwidth although imaging studies often produce large amounts of data, 10s to 100s of Mbytes, and query and retrieval times for studies will be affected by the available bandwidth and competing network traffic.

The following three scenarios illustrate the design capabilities of the MMIA:

- (i)** An individual who is at Site A for a radiological examination has previous radiological records at another hospital with a PACS. The radiologist at Site A requests the patient's previous images for comparison. How do historical images get to Site A? A solution is given in Figure 1.

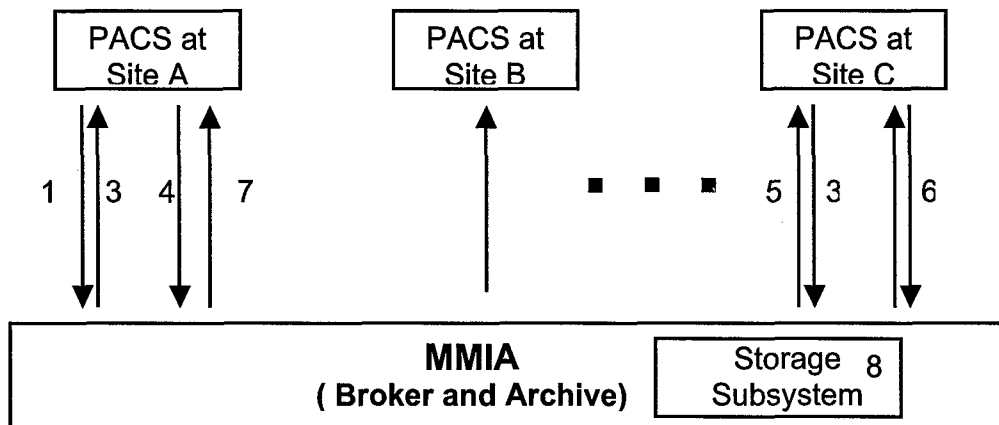


Figure 1 Broker function of the MMIA.

Procedure:

- 1: Individual at Site A. Query request sent to MMIA.
- 2: MMIA queries data from Site B, Site C, ..., until data is located at Site C.
- 3: Site C returns Individual workflow list to MMIA.
- 4: Site A selects images from workflow list and submits image retrieval request to MMIA.
- 5: MMIA submits image retrieve request to Site C.
- 6: MMIA receives images from Site C.
- 7: MMIA transmits images to Site A.
- 8: Image can optionally be stored in MMIA for future retrieval or distribution.

- (ii) An individual at Site B is temporarily relocated to Site A, which has a small clinic but no radiologist. An examination is performed with a digital radiographic unit. Images are transmitted back to Site B for reading where the radiologist decides that the current and historical images should go to Site C for an expert opinion. How should all relevant images be delivered to Site C for optimal care of this patient? A solution is given in Figure 2.

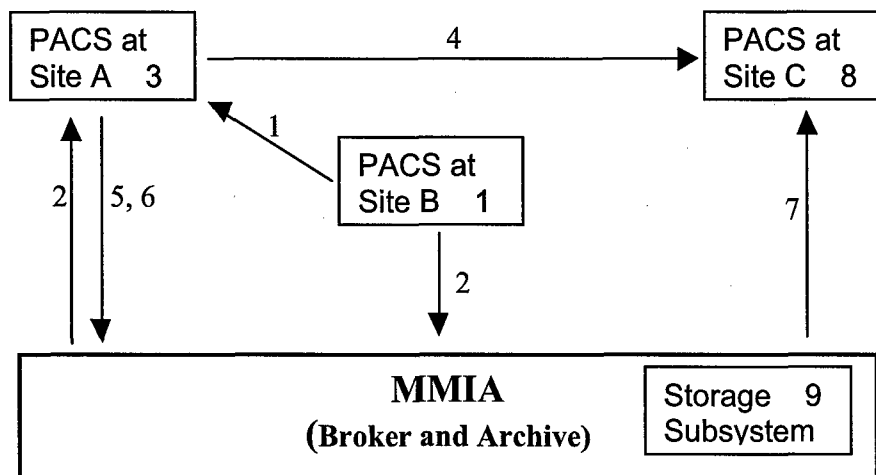


Figure 2 Broker and archive functions of the MMIA

Procedure:

- 1: Individual at Site B is transferred to Site A with a small clinic.
- 2: Site A uses the MMIA to obtain patient's images from Site B.
- 3: X-ray examination is done at Site A.
- 4: Individual is sent to Site C for further follow-up.
- 5: Site A submits Store-and-Forward request to MMIA.
- 6: MMIA receives and stores the images.
- 7: Site C requests images from MMIA.
- 8: Images from Sites A and B available at Site C.
- 9: Images from Sites A, B and C can be stored on MMIA.

- (iii) A year later, the Individual is relocated to a new site. How should we deliver the Individual's images to the new site expeditiously?

Figures 1 and 2 illustrate how the broker function of MMIA facilitates the image study transfers in scenarios 1 and 2, respectively. The third scenario considers how studies can be more easily and more rapidly accessed should the patient require healthcare in the future. If the patient's image study data is stored in the MMIA, requesting sites can obtain the data from a single transfer from the MMIA rather than waiting on an additional transfer from the originating institution to the MMIA.

Since the MMIA is the lynchpin in connecting institutions, it was designed using high availability concepts with mirrored architecture and clustering software. In the event of failure of any hardware component, the MMIA should be capable of rapid recovery and critically, without any loss in the integrity of the database. The hardware design is more fully described below.

2) Establish an inter-hospital connection (Task 2).

Connections were established between our laboratory, located within the University of California San Francisco (UCSF) complex and two other affiliated hospitals, namely Mt Zion Hospital (MZH) and San Francisco General Hospital (SFGH). All of these institutions are connected to a Synchronous Optical Network (SONET) ring, which is a redundant, fault tolerant connection with an OC-12 bandwidth (622 Mbit/sec) backbone. Other medical institutions within San Francisco also connect to this ring so the possibility for bandwidth contention does exist. The limiting bandwidth between institutions is determined by the downlinks from the SONET ring to the local routers and was 155 Mbit/sec (OC-3). Router-to-router communication across the network is done using ATM communication. Although not critical to this investigation, ATM communication links are being phased out in favor of faster and more widespread Gigabit/sec network links.

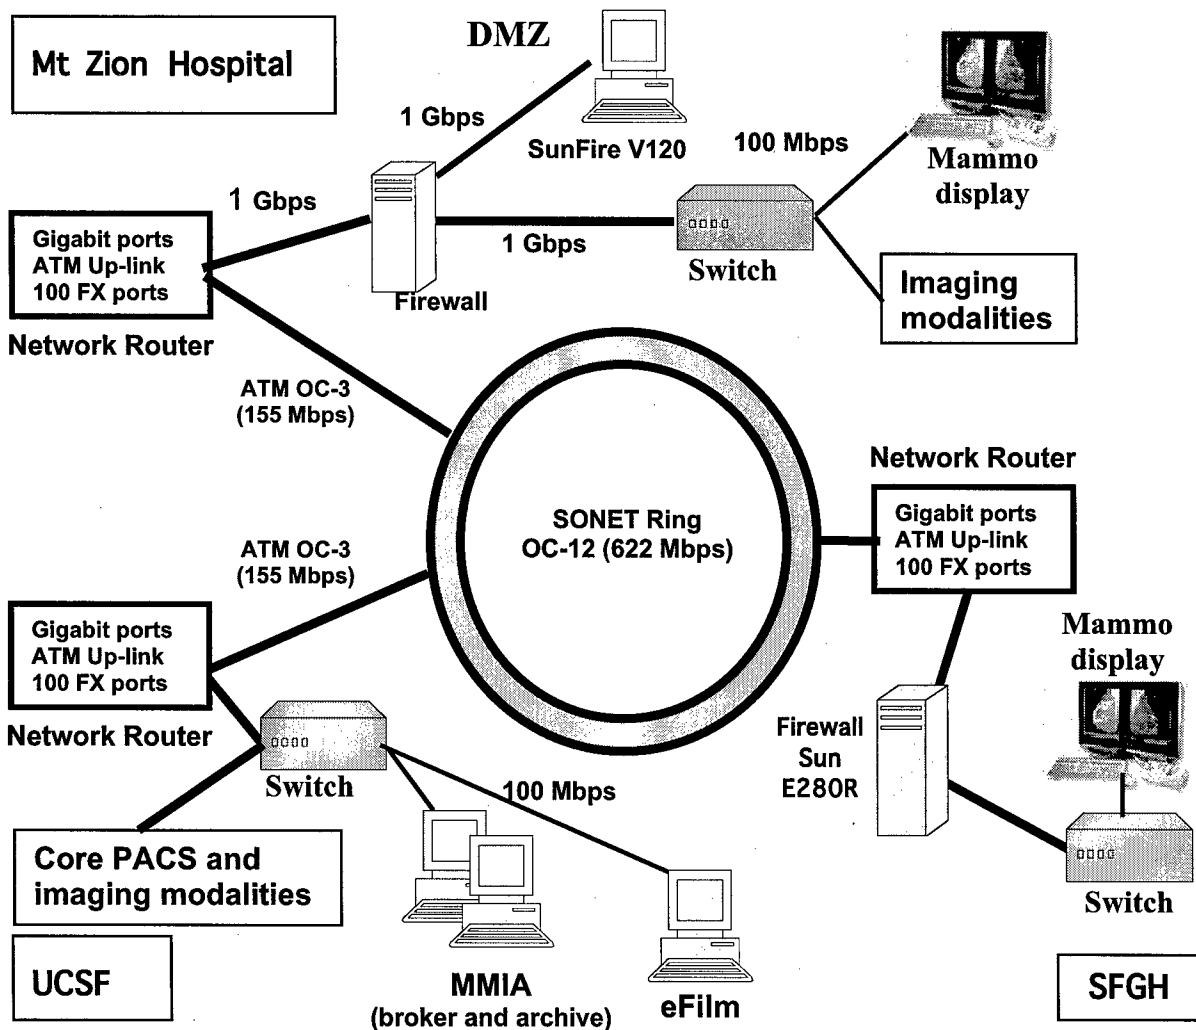


Figure 3 Hospital interconnections.

Figure 3 shows the implemented inter-hospital connections. The MMIA can connect to PACS devices at all 3 locations. At UCSF, in addition to the clinical PACS, an eFilm server (Merge Technologies Inc) was used to simulate a PACS. At SFGH, a digital mammographic display (GE Medical Systems) was used to simulate a PACS database. A firewall was implemented at SFGH to enable encryption of data between SFGH and other locations on the SONET ring. UCSF also has a firewall, which is not shown in the figure.

3) Design and implementation of the MMIA testbed (Task 3).

a) Hardware configuration (Tasks 3 a, b)

The MMIA system design was based on hardware redundancy and clustering technology in order to provide high-availability data services in a networked PACS operating environment. Figure 1 shows the configuration of the implemented MMIA system.

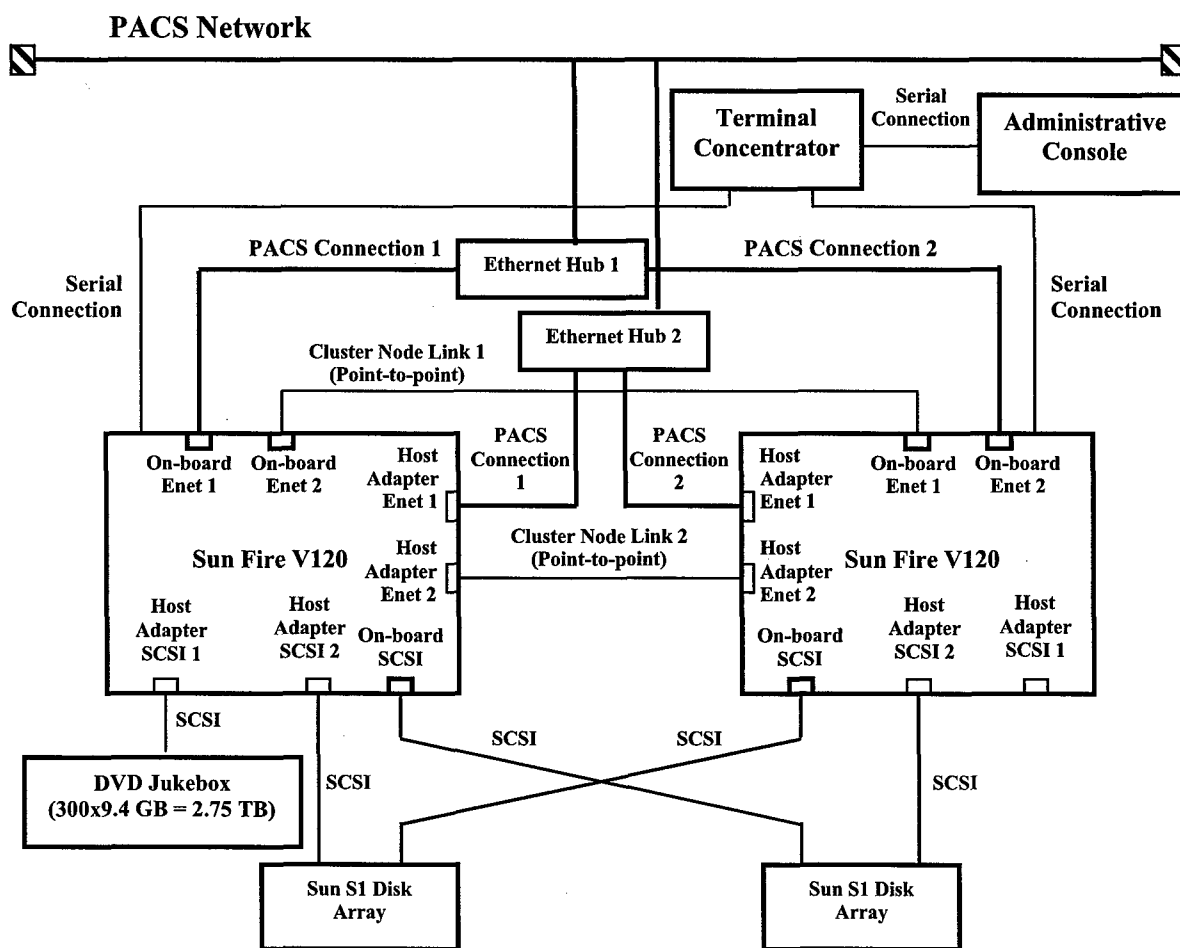


Figure 4 Meta-Medical Image Archive two-node cluster.

The high level of hardware redundancy in the MMIA is apparent in Figure 4. The MMIA contains:

- Two Sun Fire V120 systems as the primary and secondary MMIA servers
- Two Sun SCSI-based S1 disk arrays as a mirrored storage for MMIA
- A dedicated terminal concentrator (TC) connected to the V120 systems
- Dual Ethernets for the private cluster transport (inter-server heartbeat link)
- Dual Ethernets to connect the MMIA cluster to the UCSF clinical PACS network

Each CPU (Sun V120s) is connected to dual disk arrays (Sun S1s) with mirrored software. Each CPU has dual, 100BaseT Ethernet connections to the PACS network that includes the SONET ring link to remote PACS. The MMIA connects to the PACS network through separate switches to eliminate a single-point of failure. The CPUs have dual Ethernet channels for the inter-server private cluster transport, which provides the heartbeat communication required for failure detection. The Sun V120s have 2 ethernet ports and a single SCSI port. As more ports were needed, a host adapter card was added to each CPU to provide 2 additional SCSI and 2 additional ethernet ports. A Unix-based operating system, Sun Solaris 9, is implemented on these computers.

The RAID subsystem of the MMIA contains two identical RAID devices that are configured as mirrored storage. The Solaris Volume Manager (SVM), a facility embedded in the Sun/Solaris 9 operating system, creates a duplicate copy of the data so that each RAID contains identical data. The SVM software presents a virtual storage location to the MMIA application software. Thus, all write operations are duplicated while read operations come from one of the underlying mirrored RAID devices.

The mirrored RAID configuration in MMIA provides maximum data availability, which requires no failover mechanism in the event of RAID device failure. Should a RAID fail, the resynchronization process of the SVM software allows data to be automatically copied from one RAID to the other RAID. While resynchronization takes place, the RAID subsystem remains readable and writable. Thus the MMIA is immune to device failure, system crashes, or when a RAID has been taken offline and then brought back online.

A 300-platter DVD jukebox is implemented as a long-term study archive but is not been mirrored. As configured, management of the DVD jukebox prevents dual-hosting. The need for mirroring this archive was considered since failure of the CPU controlling the DVD jukebox would prevent retrieval of studies from this archive. However, the all study data within this archive would still exist on the PACS from which the data was originally retrieved and a user seeking to retrieve a study stored on the DVD archive can still obtain the study by a query and retrieve from the PACS where the study originated since this functionality of the MMIA would continue to function.

The DVD jukebox can store 9.4 Gigabytes (GB) per platter yielding a total raw storage volume of 2.8 Terabytes. Images are written to the DVD-RAM media using the industry-standard UDF file format providing transportability of the media, which consequently enables the archived images to be accessible to by any standard computer systems regardless hardware platform and operating system. However, many alternative archival storage devices exist, including digital tape and magneto-optical disks, and with the continued reduction in the cost of

RAID technology, many PACS now use ATA disk arrays. The advantage of this spinning disk storage is much faster retrieval times for data in comparison to jukebox or tape technologies.

b) Software configuration (Task 3 c, d, e)

DICOM application software to support the MMIA's image communication and storage applications uses the Mallinckrodt Institute of Radiology's CTN 3.0 utility libraries and is installed on both CPUs. Standard DICOM C-STORE, C-FIND, and C-MOVE Service Object Pair (SOP) Class services have been implemented. This software uses C programming and an appropriate compiler was purchased and installed. Table 1 lists these operations and their function. The MMIA supports both server and client DICOM processes.

SOP Class	Function
C-STORE SCP	Accepts images from remote C-STORE SCUs
C-STORE SCU	Transmits images to remote C-STORE SCPs
C-FIND SCP	Accepts and processes DICOM query requests from C-FIND SCUs
C-FIND SCU	Submits DICOM queries to C-FIND SCPs
C-MOVE SCP	Accepts and processes DICOM retrieve requests from C-MOVE SCUs
C-MOVE SCU	Submits C-MOVE requests to C-MOVE SCPs

SCP: Service Class Provider (Server Process)

SCU: Server Class User (Client Process)

Table 1 Services supported by the MMIA

To implement the MMIA function of brokering transactions between different PACS, software was implemented that allows a user to initiate a query from a local PACS for patient studies from other PACS. This functionality is achieved by having the user query the MMIA, which in turn sends the query to other PACS connected to the MMIA. This activity is shown in Figure 5 and is a DICOM C-Find request. The MMIA receives the response(s) from the remote PACS, creates a study list and returns this list to the user. The user can then request retrieval of selected images to the workstation. This retrieval request initiates a DICOM C-move request from the remote PACS archive(s) and the selected images are moved to the user's workstation. The user can also store studies to the archive of the MMIA. The activities of the MMIA are illustrated in Figure 5. The key aspect of the MMIA is that the user interacts with a single entity and not with other remote PACS. The MMIA relays user queries and user requests for studies to the other PACS connected to the MMIA.

The MMIA also queries its own local archive for studies. To manage this archive, Oracle 8i (Oracle Corporation) was installed and implemented. The Oracle database manages patient demographic data, examination history, and study and series information of studies stored on the mirrored disk arrays and on the DVD jukebox and incorporates the PACS applications

that support query and retrieval of patient examinations. These query and retrieve operations occur at three hierarchical levels via standard DICOM communication protocols:

- PATIENT Level
- STUDY Level
- SERIES Level

The implemented database software provided a two-tier hierarchical storage management (HSM) to the images that MMIA receives: A redundant RAID subsystem is used as cache storage and provides near instantaneous image access, and a DVD-RAM jukebox used for long-term image archiving (Figure 4). The data flow of the HSM software is:

- (a) Images arrive in MMIA and are stored in the RAID system and queued for permanent archiving.
- (b) Queued images in the RAID system are automatically archived to the DVD-RAM jukebox.
- (c) Archived images remain in the RAID system for fast retrieval, and are automatically purged on first-in-first-out basis.

The HSM software allows MMIA to hold the most recently retrieved images in its RAID system so that these can be accessed as quickly without having to retrieve images from the relatively low-speed DVD jukebox. The mirrored RAID system is kept as full as possible to maximize the number of images that can be retrieved quickly. However, retrieval of images from the DVD archive of the MMIA is generally faster than retrieving from a remote PACS.

- c) Implementation and testing of failover mechanism (Task 3 f, g)

Since the intent of the MMIA is to be the principal mechanism by which patient imaging data is exchanged between multiple PACS, it was designed as a high availability hardware cluster. A two-node cluster (Figure 4) was configured in an Active/Standby mode using Sun Cluster 3.1 software. In this mode, one node, the primary MMIA server, is online, carrying the full workload for the MMIA system. The other node, the secondary MMIA server, is in a hot standby mode. If the primary MMIA server fails, the secondary MMIA server automatically comes online and takes over the full workload for the MMIA system. Two modes of failure and recovery mechanisms are described below.

Cluster Node Failure

The cluster runs a kernel-resident process called the cluster membership monitor (CMM) on each node, which can detect major cluster status changes such as a node failure, a node taken off line for maintenance or loss of communication between nodes. The CMM processes communicate with one another across the private cluster interconnect network interfaces by sending regular heartbeat messages. If the heartbeat from one node is not detected within a defined time-out period, it is considered as having failed and a cluster reconfiguration is initiated (e.g., secondary node becomes primary node).

The sequence of events when a primary node fails over to a secondary node is described below:

1. The primary node is taken offline via the cluster membership negotiation process.
2. The Oracle database is un-mounted from the primary node.
3. RAID devices are un-mounted from the primary node.
4. The secondary node takes over the public (PACS) network's IP address.
5. The RAID devices are mounted to the secondary node.
6. The Oracle database is mounted to the secondary node.
7. The Oracle DBMS instances are brought up on the secondary node.
8. PACS processes on the secondary node accept data communications from the PACS network.

To achieve the last step, a duplicate PACS application is kept running on the secondary node. To the rest of the PACS network, the MMIA appears as a single IP address. The secondary node assumes this address as part of the CMM failover so there is always only a single network address for data communication.

Network Failure

In the cluster environment, both the public (PACS) network interface and the cluster transport interface on each node are dual configured (e.g., each interface is equipped with dual network host adapters). Therefore, a total of four (4) Ethernet host adapters were installed in each MMIA cluster node (Figure 4). Both the public network and the cluster transport interfaces are monitored for potential failures.

A public network management (PNM) daemon runs on each cluster node and monitors the functionality of the PACS connection. If a network connection failure is detected, the PNM automatically switches the connection to the backup. The failure and failover are transparent to the PACS applications running on each node. The cluster transport interfaces are also monitored on each node by the CMM process. If an active transport interface on one node is determined to be inoperative, both nodes route interconnect traffic (heartbeat messages) to the alternate transport interface. Again this failure is transparent to the MMIA application software.

The failover functionality of the MMIA clustering architecture was tested by physically disconnecting components from the cluster including the interconnection cable, a RAID, and network interfaces. Each of the following was tested and recovery time measured:

- Node failover
- Node network interface failover and failback
- PACS network failover
- RAID device failover
- Oracle DBMS and PACS database failover
- PACS applications failover

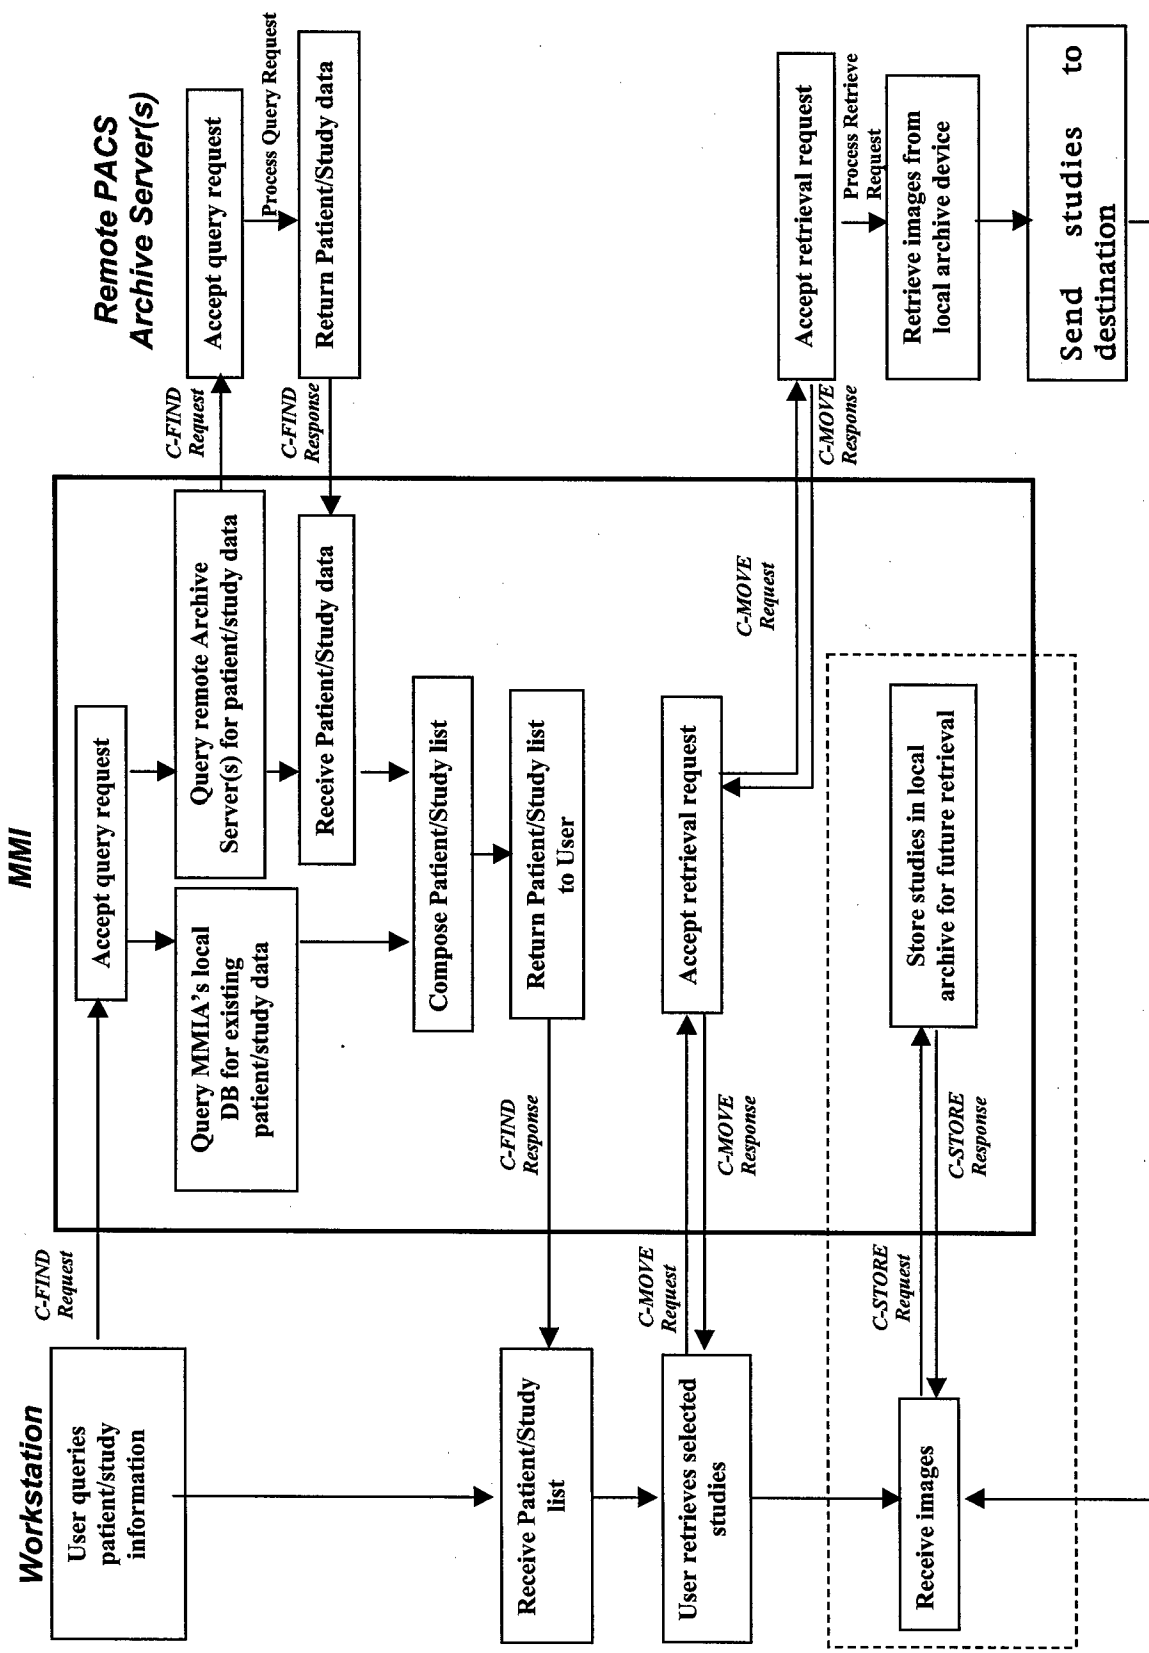


Figure 5 Query, Retrieve and Store Functions of the MMIA

Results from these tests are shown in Table 2. It can be seen from this table that initial detection of component failure or loss of network connection occurs very quickly, within seconds of the event. Failback and recover take longer depending on the problem. Worst case is failure of a node requiring the takeover of the logical name by the other node, failover of the Oracle DBMS, and activation of the application software, which requires a total time of approximately 2-3 minutes.

In the event of a RAID failure, transition to the other RAID is seamless with no loss of time. As C-Store operations continue to transfer data to the MMIA for archiving, the RAID's are no longer mirrored and a resynchronization of the RAID's is necessary when the second RAID is brought back on line. This process, which is a component of the Solaris Volume Manager software, copies data from one sub-mirror to the other to re-establish the mirrored data sets. Depending on how long the second RAID has been offline and how much data has come into the MMIA, the resynchronization process can take many hours. However, this process occurs entirely in background and does not impact the functionality of the MMIA.

Test Item	Result	Required Time
Node Failure Detection	Successful	0 – 1 sec.
Logical Node Name Takeover	Successful	28 – 32 sec.
Node Network Interface Failure Detection and Failover	Successful	1 – 3 sec.
Node Network Interface Failback	Successful	4- 6 sec.
PACS Network Connection Failure Detection and Failover	Successful	3 - 6 sec.
PACS Network Connection Failback	Successful	8 - 10 sec.
Oracle DBMS Failover	Successful	140 – 150 sec.
RAID Failover	Not applicable to mirrored storage	No delay in mirrored configuration
Data Synchronization on Mirrored RAID's	Successful	Near instantaneous
Data Re-synchronization on RAID's after one RAID fails	Successful	90 – 120 min.* *Based on data stored on 2x36-GB RAID.
PACS Applications Failover	Successful	4 sec.* *After completion of Oracle DBMS failover

Table 2 Test results of failover and recovery of the MMIA.

4) Establish connections to the MMIA (Task 4)

The connections of the MMIA to multiple PACS was described in (2) above and shown in Figure 3. In addition to the connection to the core PACS at UCSF, and eFilm server (Merge Technologies) was implemented for testing purposes. This device has a local database and supports DICOM C-FIND, C-Move and C-Store SCU and SCP services (ref). Thus this eFilm server looks to the network as a full PACS entity.

We did not connect the PACS network to any digital mammographic system although this had been planned. No digital mammographic unit was implemented within the UCSF system during the grant period. However, mammographic images were available from a full field digital mammographic (FFDM) unit, a GE Medical Systems Senographe 2000D, that was installed in a mobile van. During patient imaging operation, the van was not connected to the network; rather in the evenings it drove to a homeport and connected to the UCSF PACS port for a time sufficient to transfer images to the PACS. These images were viewed on the mammographic display located at MZ hospital. The mammographic display at San Francisco General Hospital could then query for these images.

5) Implement information protection infrastructure (Task 5)

All data transfers between any PACS and the MMIA must be secure. This security could be achieved by communicating through private networks but our prototype system was implemented to use either public or private networks. As shown in Figure 3, the MMIA, the UCSF PACS and the eFilm server all reside on the same network behind a common firewall. However, the mammographic display station at SFGH connects via the SNET ring, which is considered a public network. To establish the secure transfer of data across this link, a commercial encryption software product made by Check Point Software Technologies was implemented on the firewall at SFGH and at UCSF. Thus transmissions across the SNET ring could be encrypted using the Check Point software, which uses the Advanced Encryption Standard (AES), 128-bit encryption algorithm and private key encryption. MD5, a data integrity algorithm, is incorporated into the Check Point software.

While encryption is necessary over any public network, the effective transmission rate of the data will be slowed because of the time required to encrypt and de-encrypt the data. The time required for these activities depends on the power of the CPUs and also on the characteristics of the images being encrypted. To determine the impact of encryption, we defined two test data sets, one with FFDM images and one with MR images (GE Medical Systems). Table 3 shows some features of each of these images sets. It can be seen that the principal difference between these two types of images is the bit volume per image with the MR images almost a factor of 20 smaller in volume.

Measurements were made of the of disk-to-disk transfer times between the two firewall computers with and without the Checkpoint encryption software enabled. Results are shown in Table 4. Images were transferred using the DICOM protocol. It is apparent that while encryption is critical for data integrity and confidentiality, it adds significant overhead, reducing the effective transfer rate. In practice the observed reduction will be both hardware and software dependent, but any implementation will reduce throughput. Furthermore, if the firewalls are in

heavy use because of other network traffic, the impact of implementing software encryption will be greater.

Alternatively, the Checkpoint software could be implemented within the MMIA. We did not choose this alternative as this could slow the broker functions of the MMIA. It was felt that the impact of encryption could be minimized by use of a dedicated firewall so that the encryption process does not slow activities and traffic behind the firewall.

Equipment	Pixel size, mm	Matrix size	Pixel bit depth	Image size, Mbytes	Number of Images	Study Bit volume, Mbytes
GE FFDM	0.01	1914 x 2294	14	~ 8.5	24	202
GE MRI	~ 0.5	512 x 512	16	~ 0.5	696	354

Table 3 Test data sets for encryption testing.

Image type	Data transfer rate reduction
GE FFDM	53%
GE MRI	26%

Table 4 Reduction in disk-to-disk transfer time by encryption.

6) Evaluation of the MMIA (Task 6, 7)

Initial tests of the broker function of the MMIA were conducted using a commercial DICOM application, eFilm, to query a remote PACS archive (the UCSF PACS) and retrieve selected studies from this archive. Queries for studies were done with patient, series and image requests to verify the hierarchical implementation of the DICOM functions. The MMIA relayed service requests to the secondary PACS and the responses of the secondary PACS back to the workstation initiating the requests. MR, CR, CT and digital mammographic studies were retrieved via broker transactions.

Timing studies were done to compare queries and retrieval times when the secondary PACS was queried directly compared to use of the broker. These tests did not require the data cross a firewall boundary and no encryption of the data occurred. No difference could be detected between these methods. However, in practice outside the controlled environment of our testbed, many factors influence such timing tests including bandwidth and competing traffic on the networks. Furthermore, such times are highly dependent on the database size and priority settings of the responding PACS. That is, the response of commercial PACS to remote DICOM queries is often assigned a lower priority than internal activities including image acquisition from

modalities, distribution of studies to workstations and archiving and retrieval of images from the local database. At busy times, PACS activities can be significant and outside queries could be delayed.

The MMIA was not used extensively in a clinical setting in part because a digital mammographic unit was not installed in any fixed setting. However the broker was used to move images from the UCSF archive to the mammographic display station at SFGH for review. This display workstation could also display MR studies and these were also accessed via the broker. In these applications, encryption did occur between firewalls so that communication across the public SONET ring was protected. We also compared the time required to move studies between UCSF and SFGH to those obtained when the same studies were within our laboratory. That is, we wanted to determine the impact of retrieving studies over a wide area network.

As shown in Figure 3, the wide area network (WAN) that extended to SFGH was a high bandwidth network, limited by the 100 Mbit/sec Ethernet ports of the MMIA and similar ports of the other DICOM components. A closed local area network (LAN) was established within our laboratory with the MMIA and eFilm components attached with 100 Mbit/sec links. Disk-to-disk transfer times were measured for the test data sets shown in Table 3. These measurements were then repeated across the WAN and compared to the LAN. All measurements used the DICOM C-move protocol. Measurements were made without encryption of the data. Results are shown in table 5.

Image type	Bandwidth Utilization (% of maximum)		Reduction in transfer rate
	LAN	WAN	
GE FFDM	45.8%	25.5%	44.2%
GE MR	18.4%	14.9%	19.1%

Table 5 Local versus wide area data transfers.

The effective transfer rates depended on the data being sent as was observed with encryption. The use of a wide area network did reduce the effective transfer rate as expected since transmissions across the WAN competed with other network traffic. However, the effective transfer rates were still in excess of 10 Mbit/sec. The average transfer time across the WAN per image was 2.7 sec for the FFDM image and 0.26 sec for an MR image. Practical use of the MMIA is dependent on high-speed network links to remote PACS to minimize image retrieval times.

Research Accomplishments

- A high availability Meta-Medical Image Archive has been designed and implemented that is capable of brokering transactions between Picture Archiving and Communication Systems at multiple institutions. The MMIA features state-of-the-art failover architecture including dual CPUs, mirrored disk arrays and redundant connections to a PACS network.
- A broker function has been implemented on the MMIA that allows query and retrieval requests, initiated by a user at a workstation of a local PACS, to be relayed to a third-party image archive (e.g., a remote PACS). The MMIA relays all query, response and retrieve requests allowing the user to communicate with a single entity while interrogating and retrieving from multiple locations.
- The MMIA broker function allows multi-vendor image archive systems to be virtually interconnected via MMIA. The varied DICOM information models (e.g., Patient Root query and Study Root query) adopted by individual PACS systems make the integration of data from multiple PACS difficult. The MMIA hierarchical database scheme enables responses to queries to be presented to a user in a single coherent view.
- The MMIA has its own database and storage capabilities allowing retrieved studies to be stored within the MMIA for future retrievals. Both short-term storage, in the form of mirrored RAID arrays, and a long-term archive, a DVD-RAM jukebox, have been implemented.
- The broker activities of the MMIA have been tested using multiple modality images including MR, CT, CR and digital mammograms, representing studies having a large data volume per study and per image.
- The MMIA is designed as a high availability cluster. Failover, failback and database recovery (synchronization) following failure of component hardware, network connectivity and software have been successfully tested.
- The MMIA has been shown to function across a wide area network. The time taken to retrieve images from remote PACS will depend on network bandwidth.
- The impact of software encryption technology on transfer times of images between a remote PACS and the MMIA has been evaluated.

Reportable Outcomes

- 1) Poster presentation at Peer Reviewed Medical Research Program (PRMRP) Investigators Meeting held in Puerto Rico, April 26-28, 2004. (Appendix 1)
- 2) Abstract submission to the Scientific Assembly of the Radiological Society of North America November 28 – December 3, 2004. Received alternate status. (Appendix 2).

Conclusions

The growth and expansion of Picture Archiving and Communication Systems in radiology has the potential to significantly improve access to imaging studies of an individual taken at multiple institutions. To replace film transport, some institutions are burning studies onto CDs and the patient can hand carry these between institutions. But the ability to remotely query and retrieve studies from multiple PACS has not yet been implemented extensively. Thus a functional broker that could identify image data of a given individual and manage its transfer to a user at a remote workstation is of significant value.

Transactions between remote PACS including query and retrieve functions can be accomplished using a brokering server. Relevant imaging studies of an individual, acquired at different institutions, can be retrieved to the local workstation from which the request was initiated without directly querying the remote PACS. The MMIA can be designed with fault-tolerant architecture to assure high reliability. This high-availability architecture assures both data integrity and near constant uptime.

Appendix 1

Background

Many medical facilities now acquire, store, distribute and display images electronically. Deployment of Picture Archiving and Communication Systems (PACS) in hospital settings has grown at a double-digit rate in recent years because of recognition of the efficiencies PACS brings to an institution and because imaging modalities are now designed to take advantage PACS, producing vast numbers of images per study that preclude filming. A key element of PACS is the standardized communication protocol (DICOM) that not only allows imaging equipment of all makes and models to be coupled to a common network, but also allows users with appropriate permissions to query and retrieve studies from a remote DICOM server to any DICOM compliant device or system. Thus, if a patient accesses care at multiple institutions, imaging studies from both institutions can be retrieved and integrated. However, this process must protect patient data, be efficient and reliable in its operation and be implemented without disrupting local PACS operation.

To address problems with this integration that stem from different information infrastructures at different institutions, a meta-manager can be used as a systems broker, interconnecting different databases using DICOM protocols. The purpose of this project was to develop a Meta Medical Image Archive (MMIA) test-bed that interconnects multiple PACS and addresses issues of reliability, security, and efficiency.

Methods

Design and implementation of a Meta Medical Image Archive (MMIA) test-bed

1. MMIA Cluster Design

To assure high reliability of the MMIA, a redundant CPU architecture using failover software has been implemented. This dual-server MMIA cluster provides high-availability services with failover capability and is shown in Figure 1. The high level of redundancy is apparent. Each CPU (Sun V120s) is connected to dual disk arrays (Sun S1s) with mirrored software. Each CPU has dual, 100BaseT Ethernet connections to a PACS network to which remote PACS are connected. In our testbed configuration, three separate PACS have been simulated using DICOM compliant display, archive and server systems (Merge Technologies eFilm). The PACS network resides behind a firewall and in a clinical setting; the remote backs would be similarly protected. Thus the connection between PACS can occur by any digital link, public or private. Clearly, the bandwidth of the connecting links and competitive traffic will influence data transfer rates. The MMIA is connected to the PACS network through separate switches to eliminate a single-point of failure. In addition, the CPUs have dual Ethernet channels for the inter-server private cluster transport, which provides the heartbeat communication required for failure detection.

The long-term archive, a DVD jukebox, has not been mirrored because management of the DVD prevents dual-hosting as currently configured. However, this archive is redundant storage as the studies are also archived on the PACS from which the data was initially obtained. In the event of

failure of the CPU controlling the DVD, the alternate CPUs broker functions will still be functional. Thus a user seeking to retrieve a study that is stored on the DVD archive can still obtain that study by a query and retrieve from the PACS where the study originated. The consequence would be some time penalty as it would be faster to satisfy queries from a remote PACS using the broker's archive since PACS frequently have a priority setting for processing transactions with remote entities and if that priority is low and/or at peak times of the day, response times may be slowed due to slow response of the local PACS.

The MMIA cluster has a virtual IP address seen to the network and has a local database (Oracle), which is shared by both nodes but only a single server can access the database at one time. This primary node performs all functions of the MMIA including brokering transactions, the Oracle service process and data storage. While no load sharing occurs between the two nodes, all application software and the Oracle service process is running on both nodes at all times. The Oracle database and all storage objects are mirrored on the RAIDs. In the event of failure of the primary node, the following operations occur:

- Node failure detection
- Mounting of the Oracle storage devices on the secondary node
- Activation of the Oracle service process on the secondary server
- Transfer of logical host name (IP address)
- Application failover

Node failure and the time required for recovery have been tested. Failures of a RAID or of a network interface are seamless and have also been tested.

2. MMIA Broker Software

The DICOM application software to support MMIA's image communication and storage applications uses the Mallinckrodt Institute of Radiology's CTN 3.0 utility libraries and is installed on both CPUs. Standard DICOM C-STORE, C-FIND, and C-MOVE SOP Class services have been implemented.

Brokering software has been implemented that allows a DICOM query for patient studies, initiated by a user at a remote PACS, to be relayed to other archives. The MMIA receives the response from the remote system, creates a study list and returns this list to the user. The user can then request retrieval of selected studies to the workstation. This retrieval request initiates a DICOM C-move from the remote archive and the selected studies are moved to the user's workstation. The user can also store studies to the archive of the MMIA. These activities are illustrated in Figure 2.

3. Security

Image data transfers between PACS often use public networks and must be secure. The MMIA must itself be safe from attacks and security breaches. The MMIA has been deployed behind our Departments' firewall and all PACS need to be similarly isolated. We have tested the MMIA by

querying a DICOM compliant server located behind another firewall at a companion institution. A secure software program that is considered the industry standard, Checkpoint, was used to encrypt all transfers between firewalls. Encryption does impose some penalty on the transfer time but when implemented between firewalls, does not affect broker functions.

Results

Tests have been conducted using a commercial DICOM application, eFilm, to query a remote archive and retrieve selected studies from this archive with all transactions brokered through the MMIA. That is, the MMIA server relayed service requests to the secondary PACS and the responses of the secondary PACS back to the system initiating the requests. MR, CR, CT and digital mammographic studies were retrieved via broker transactions. Timing studies were done to compare queries and retrieval times when the secondary PACS was queried directly compared to use of the broker. No difference could be detected between these methods. However, in practice outside the controlled environment of our testbed, many factors influence such timing tests including bandwidth and competing traffic on the networks. Furthermore, such times are highly dependent on the database size and priority settings of the responding PACS.

Table 1 shows the response of the MMIA to component failures. Should the primary node fail, recovery occurs in a matter of minutes without operator intervention. The time required is the sum of detection, Oracle switchover, logical IP transfer, and application failover.

Component/item	Required Time
Node failure detection	<< 1 sec
Oracle switchover	110-120 sec
Logical host transfer	28-32 sec
Broker application failover	12-15 sec
Network interface card failure	<< 1 sec
RAID failure detection	<< 1 sec

Table 1 Failover and component failure tests

Conclusions

The growth and expansion of Picture Archiving and Communication Systems in radiology has the potential to significantly improve access to imaging studies of an individual taken at multiple institutions. To replace film transport, some institutions are burning studies onto CDs and the patient can hand carry these between institutions. But the ability to remotely query and retrieve studies from multiple PACS has not yet been implemented extensively. Thus a functional broker that could identify image data of a given individual and manage its transfer to a user at a remote workstation is of significant value.

Transactions between remote PACS including query and retrieve functions can be accomplished using a brokering server. Relevant imaging studies of an individual, acquired at different institutions, can be retrieved to the local workstation from which the request was initiated without directly querying the remote PACS. The MMIA can be designed with fault-tolerant architecture to assure high reliability.

Appendix 2

DEVELOPMENT OF A META-MEDICAL IMAGE ARCHIVE (MMIA) FOR INTEGRATION OF IMAGE DATA FROM MULTIPLE INSTITUTIONS

Gould RG, Wong A

Department of Radiology, University of California, San Francisco, CA 94143

robert.gould@radiology.ucsf.edu

BACKGROUND/PURPOSE: Many hospitals have implemented Picture Archive and Communication Systems (PACS) for the distribution, display and archiving of diagnostic imaging studies, providing the potential to move these studies electronically between institutions. This study presents the design of a prototype MMIA that is capable of brokering and facilitating transactions between multiple institutions. The MMIA provides a mechanism for a practitioner at one institution to identify, retrieve and store for future-use relevant imaging studies of an individual that have been acquired at different institutions. The query is made to the MMIA, not directly to other PACS. **METHODS:** An MMIA testbed has been designed and implemented that is highly redundant, utilizing a state-of-the-art clustering architecture. The MMIA has been connected to multiple simulated PACS and a broker function has been implemented on the MMIA that allows a DICOM query, initiated by a user at a remote PACS, to be relayed to other PACS. The MMIA receives the response from the remote archive, creates a study list and returns this list to the user. The user can then request retrieval of selected studies to the local workstation. This retrieval request initiates a DICOM C-move, relayed through the MMIA, for the remote archive to move the selected images to the user's workstation. The MMIA has a local database and archiving capabilities. **RESULTS:** The function of the MMIA has been tested by relaying requests for studies between simulated remote PACS. CR, MR, CT and digital mammographic images were successfully retrieved via the broker function. The MMIA allowed studies to be identified and retrieved on both a study and series basis. The fail-over capability of the MMIA has been validated by the forced interruption of both CPU and network components. **CONCLUSION:** Transactions between remote PACS including query and retrieve functions can be accomplished using a brokering server. Relevant imaging studies of an individual, acquired at different institutions, can be retrieved to the local workstation from which the request was initiated without directly querying the remote PACS. The MMIA can be designed with redundant architecture to assure high reliability.

The U.S. Army Medical Research Materiel Command under DAMD17-00-1-0585 supported this work.