

Building a Better Medical Image Archive



**A White Paper
Prepared by**

**Michael J. Gray
Gray Consulting**

**P.O. Box 972
Novato CA 94948-0972**

graycons@well.com

Introduction

On October 26, 2006 the first part of a two-part article dealing with the issue of data migration from PACS to PACS was published in auntminnie.com's Letter from the Editor. The response to this article as well as the second part published in the January 4, 2007 Letter from the editor has been very favorable. A number of emails requested additional clarification and an expansion on the subject. This white paper is a response to those requests. It is an organized presentation of the various technical issues that must be overcome if we are going to achieve the development of an Enterprise Archive and thereby eliminate the requirement for data migration every time a PACS is replaced.

The various concepts, problems, and solutions described in this paper were learned over the course of several consulting projects. I thank my clients for providing me the opportunity to continue learning about the inner workings of Picture Archiving and Communication Systems. I thank the various vendors with whom we engaged during these projects for being so willing to share their ideas and technical information.

Building a Better Medical Image Archive

Making the change from PACS A to PACS B requires the migration of study data from A to B, no matter how DICOM-conformant either system is said to be. It is simply not possible to unplug the archive from PACS A and plug it into PACS B.

There are two principal reasons for this inevitable data migration. The first reason has to do with the differences in the way the two vendors create, interpret or utilize the DICOM file headers. In this case, the data migration process is required to convert the data created in PACS A into a format that can be interpreted and utilized by PACS B. The second reason has to do with the fact that most PACS vendors refuse to disclose their Directory's Data Dictionary or its Schema, the keys to understanding what patients and studies are in the PACS and the file system pointers that describe where the study data is actually stored in the attached storage solutions. So even if PACS A and PACS B were fully DICOM-conformant, and both vendors created nearly identical DICOM headers, the data would have to be migrated from A to B in order for PACS B to "learn" all of the patients and studies stored in PACS A, and then create its own file system pointers.

How big is this data migration task? Let's look at the example depicted in the Table below. A single facility performing 100,000 Radiology procedures in 2001, installs a PACS. Their annual growth in procedure counts averages 5% for all study types. In 2006 they decide to replace their first PACS with another PACS. This will require migrating approximately 22 Terabytes of study data. A large but manageable task.

	01	02	03	04	05	06	07	08	09	10	11	12
New (TB)	4.0	4.2	4.4	4.6	4.9	5.1	5.4	5.6	5.9	6.2	6.5	6.8
64 slice (TB)						4.2	4.4	9.3	9.8	10.2	10.8	11.3
DM (TB)						1.2	1.3	1.3	1.4	1.5	1.5	1.6
Cum. (TB)	4.0	8.2	12.6	17.2	22.1	32.6	43.7	59.9	77.0	94.9	114	133

Estimated data equivalents over a 12 year span starting in 2001.

In 2006, the same year that they migrate to their second PACS, the facility decides to replace the older of its two CT scanners with a 64 slice model and add two Digital Mammography units. The additional data created by those devices is indicated in the Table. In 2008, the facility replaces its second CT with another 64 slice model. If the facility decides to move on to yet another PACS in 2011, they will be looking at having to migrate an estimated 114 Terabytes of data, less whatever they have been able to purge. Assuming that they can purge every study that is over 7 years old, that would mean they could possibly eliminate the study data from 2001 through 2003 (roughly 12.6 TB), thus reducing their migration to a mere 100 TB. In six years their data migration task grew five-fold, and we did not factor in additional modalities and newer imaging technology that will undoubtedly create larger study data sets.

In reality, they will not be able to purge the pediatric studies, or studies belonging to employees, worker compensation cases, or studies belonging to patient's that have survived on-going conditions that will require further treatment. It is widely believed that the lifecycle of a growing percentage of radiology studies may reach out to 20 years. Given the probable growth in study volumes, images per study, and image matrices, the actual data migration that our sample facility faces is probably considerably larger than 100 TB in 2011, and the task may be simply overwhelming another six years later. Clearly even the most current generation of PACS is not designed to manage the entire lifecycle of the radiology study, if the lifecycle of the PACS itself spans a mere five to six years. The sobering truth is that a facility will probably have to migrate a large percentage of its study data as many as three times, as many times as it is likely to change PACS over the course of twenty years.

So the most important question related to data migration is, "What can be done to eliminate the need to migrate data every time a facility changes PACS?". Stated another way, "What is the key to the 'Never Again' strategy, where the data is migrated once, and Never Again?".

The simple answer is to reinvent how radiology study data is managed. If the PACS (as we currently know it) cannot manage the complete lifecycle of radiology study data, then the data management and the archive functionality of the PACS must be separated from the PACS and embodied in a new kind of data management system.

Naming this new kind of data management system an "Archive", doesn't do the concept justice. This new data management system would have to be "intelligent", meaning that it would have to have its own database management component that would continuously manage the study data over its full lifecycle. Since this data management system would be capable of managing Radiology, Cardiology, and any other '-ology' study data, we should think of it as a general purpose Medical Information Management System (MIMS). The MIMS would be architected to be highly reliable and massively scalable, and the hardware configuration would simply evolve through 20 years of change in storage technology.

The PACS (as we currently know it) would then "evolve" into a server subsystem that would focus primarily on supporting specialized information distribution and display applications, and there would probably be such a system for each of the medical departments that create images: Radiology, Cardiology, Pathology, etc. We should think of each of these systems as a Medical Information Distribution and Display System (MIDDS).

As is the case with most radiology PACS today, the Radiology MIDDS would focus on Distributing and Displaying Radiology Information. This system would provide one set of applications that would constitute the Diagnostic Display Station, another set that would constitute the Specialist Display Station, and another set that would constitute the Generalist Display Station. In many current generation Radiology PACS, there is only one master display application, and each class of display station is actually created by


turning various features on or off based on the display requirements of the various users, Radiologist, Surgeon, and Generalist.

The Cardiology MIDDs would provide a similar family of display applications specializing in the distribution and display of Cardiology Information. Another MIDDs will eventually emerge to focus on Pathology Information.

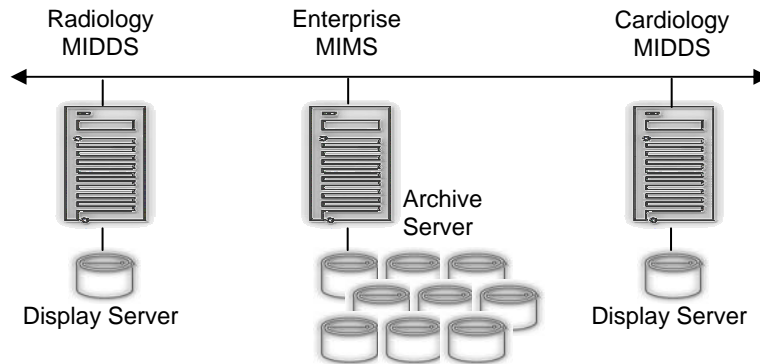
Each MIDDs would create and manage user work lists and have its own local disk cache that would be configured to store a user-defined number of months of new study data and any related priors. One to two months of study data would cover most clinical needs, but some organizations might want to keep as much as the most recent year of study data on this local cache. In any event, it would no longer be necessary for a MIDDs to have its own long-term archive component. Responsibility for long-term information management (up to 20 years) would fall to the single, shared enterprise MIMS.

It is interesting to note that many of the PACS vendors have been hinting at the very specialization I am suggesting in this paper. By RSNA 2006, most PACS vendors had begun offering a choice of storage solutions with their PACS. In many cases, EMC², HP, IBM, etc. are calling directly on the Health System and making direct proposals for storage solutions. This is essentially a signal that the PACS vendors are willing to loosen their grip on the storage solution. At the same time, the PACS vendors are admitting that their expertise, their true value-add is the display application. Modality vendors such as GE, Philips, and Siemens are integrating image processing applications developed for their imaging modalities directly into their PACS display application. The day is not far off, when a PACS will be little more than a Master Display Application delivered on a series of installation CDs.

The set of graphics that follow illustrates the suggested evolution of PACS (as we know it today) into the various subsystems that would be better equipped to support the needs of the Health System.



Conventional departmental PACS, each with its own dedicated long-term archive, would evolve into separate departmental Medical Information Distribution and Display Systems and a shared, enterprise Medical Information Management System.



By design, the MIMS would have to be compatible with each successive generation of MIDS, and each MIDS would basically become the support structure for user worklists, work flow management and display applications. These are the applications that today's PACS vendors are investing most of their R&D dollars into, and the applications which they see as the key to their product differentiation.

I apologize for creating new acronyms for these various subsystems. I simply want to point out that the familiar terms "PACS" and "Archive", wouldn't really apply in this new configuration. The conventional PACS would be subdivided into separate systems, one that would focus on information distribution and display and one that would focus on information management. However there may be less confusion, if from this point forward we use the more familiar terms, so my Medical Information Distribution and Display System (MIDS) will hereafter be referred to as a mini-PACS, and my Medical Information Management System (MIMS) will hereafter be referred to as the Enterprise Archive.

DICOM connectivity requirement

The key to building any system that is itself composed of two or more subsystems from different vendors is intersystem connectivity...how study data is exchanged between the subsystems. I would argue that DICOM and not HTML is the better technology choice for this interface. A higher percentage of PACS vendors either already support a complete suite of DICOM or will support a complete suite of DICOM in their next software release. Assigning a URL tag to every data object and converting every PACS subsystem into a web server has its merits, but evidence suggests that the DICOM standard will be more widely embraced and sooner.

An Enterprise Archive that supports the complete suite of DICOM data objects and SOP Classes would then be able to exchange data objects with any DICOM-conformant PACS.

While it is encouraging that many of the newer PACS are about to become more DICOM-conformant, what of the many PACS that are already installed? How DICOM-conformant are they? How easily could they be interfaced to this Enterprise Archive? The unfortunate answer is that many PACS are DICOM-conformant on their front end, receiving DICOM study data from the imaging modalities, but not very DICOM-conformant on their rear end. They still don't play very well with other vendor's systems.

Even in a conventional PACS, one that includes its own long-term archive, study data is exchanged between various servers that make up the system. But the developers of most currently installed PACS were free to use whatever data transfer mechanism was the easiest, the most efficient, etc. for their intra-system data transfers. They were not obligated to conform to any standard method of data exchange. Not only did these developers have no requirement to exchange study data with another vendor's subsystem, it can be argued that they saw no advantage to making it easy to interface to another vendor's system. As a consequence, most of the PACS installed today are fairly closed systems. Once the data has been acquired by one of these PACS, it is pretty much captive in that PACS.

Lack of DICOM conformance is one of the key reasons why the standalone DICOM Archive failed to catch on six years ago, and it is a key problem that must be overcome if the standalone Enterprise Archive is to succeed this time around.

In order to fully understand the new Enterprise Archive concept, it is necessary to review the various levels of DICOM conformance that would be required in order for the Enterprise Archive to exchange data with a new or even a DICOM-weak incumbent PACS.

There are four areas where this DICOM conformance would be required

- **Data Objects.** This would include the study images; associated reports; so-called meta data objects associated with the images like Patient Name, Patient Medical Record Number and other patient/study identifiers; Study Accession Number; Grayscale Softcopy Presentation States (GSPS), Key image Notes (KIN), and Presentation of Grouped Procedures (PGP).
- **Image Header Tags.** These are place holders in the image header where the meta data objects are supposed to be located. DICOM describes the location and the format for storing these objects. Public Tags are required for name, ID, etc. and Private Tags are allowed for other study characteristics such as study status, modality set-up, etc.)
- **Data Compression Syntax.** DICOM supports use of J-PEG 2000.
- **Intersystem Data Exchange Operations.** (Query/Retrieve, Find/Move), especially automated pre-fetching from a foreign archive.

Developers of many incumbent PACS would argue that if the system is comprised of components and subsystems from the same vendor, there is no real need to be fully DICOM-conformant. While that is a somewhat self-serving argument, it is also self-fulfilling, for a lack of DICOM conformance in all of these areas will virtually assure that all of the components and subsystems will have to come from the same vendor. Worse, it would be necessary for all of the components to be from the same generation of a vendor's PACS. Time and time again we have seen what happens when the vendor introduces a next-generation PACS...the previous generation components and subsystems must be completely replaced!

DICOM conformance in each of the areas outlined above would be required in order to mix components and subsystems from different vendors to create a single interoperable system. This is the epitome of the "Best-of-Breed" strategy.

As is the case today, it is unlikely that any one vendor would have the best Radiology mini-PACS, the best Cardiology mini-PACS and the best Enterprise Archive. More likely, just as is the common practice today, a Health System would probably purchase such systems at different times and from different vendors. Therefore the only way to guarantee interoperability between these systems is to make sure that each individual system is fully DICOM-conformant in all four of the areas listed above:

- Data Objects
- Image Header Tags
- Data Compression Syntax
- Intersystem Data Exchange Operations

To state the key objective in simple terms, each individual mini-PACS and the Enterprise Archive must be able to bi-directionally exchange all types of data objects using DICOM SOP Classes.

This is not the time or the place to undertake a detailed review of DICOM, but a few clarifying examples might be helpful.

- Data Objects

Most installed PACS allow the radiologist to create a Presentation State (save the window/level set-up and an annotation or graphic overlay on an image) and a Key Image Note (set a key image flag and type a text-based note related to a key image), but they do not treat these data objects as DICOM Objects. They are treated instead as proprietary objects and stored not in the image header but in the system Directory. PACS that behave in this manner not only cannot pass these proprietary objects on to the Enterprise Archive and any other PACS, they cannot receive and display GSPS or KIN created by any other system. These important work products are simply trapped in the PACS.

- Image Header Tags

The DICOM standard carefully specifies where in the image header specific information (meta data) must be located and the precise format of that information. These locations are referred to as Tags. Key information such as Patient Name, MRN, Study UID, Study Name, Referring Physician, etc. must be stored, as described, in Public Tags. Public Tags (those containing key identifying information) are generally supported by the PACS vendors. But there are many Private Tags that a vendor is free to use to store important information like study status (unread/read), or special overlays, notes, etc. A system could be DICOM-conformant, yet unable to share their Private Tag information with another system. For the same reason, the same system might be unable to interpret similar information buried in the private tags written by another system. The ability to recognize the study status (unread/read) would impact a system's ability to differentiate between a new study and a relevant prior, thus impacting that system's pre-fetch operations. One of the functions of a data migration application is to make sure that the new PACS recognizes the data being migrated from the old PACS as "read" data and not "unread" data, so the migration application has to know how to make the appropriate tag and the information stored in that Tag recognizable to the new PACS. Extracting information stored in Private Tags by one PACS and placing it in the appropriate Tag so that it can be recognized by another PACS is what differentiates Data Migration services, and requires a great deal of field experience with the various PACS.

- Data Compression Syntax

Data compression is a still a very useful tool. Despite faster networks and cheaper disk drives, you can never have enough bandwidth or enough storage capacity. But some compression schemes are better than others. Unfortunately a number of PACS vendors see another advantage to data compression...vendor lock. With today's technology, there is no good argument to the use of proprietary compression schemes. DICOM now embraces J-PEG 2000, both the loss-less and lossy variations. Systems that only support proprietary schemes must first reconstruct their compressed data before it can be transferred to another system. This requires both CPU cycles and bandwidth, and you can never have enough of either.

In order for mini-PACS and Enterprise Archives to effectively work together, all data objects must be stored as DICOM objects in the proper format, in their proper Tag in the DICOM header and compressed according to a DICOM-conformant compression scheme. Until this is the case, there will continue to be the need to extract non-DICOM objects from Directories and convert them to DICOM objects, the need to extract information located in Private tags and place it in Public or Private Tags that can be read by the other system, and the need to waste bandwidth transferring uncompressed data files between systems. These extraction and conversion routines will continue to be required every time an image is exchanged between two systems that are not 100% DICOM conformant, or if one or both systems have buried important information in Private Tags. This data extraction and conversion process is commonly referred to as "morphing".

It is interesting to note that some of the vendors have developed considerable skill in morphing their header to look like that of their competitor's, and their competitor's header to look like their header. These skills are required if a vendor wants to participate in the PACS replacement market. Rather than support both the new PACS and the old PACS, most users expect their new PACS vendor to migrate all of the study data from the old PACS to their new PACS, so their old PACS can be decommissioned. This data migration from old to new PACS frequently requires header "morphing".

- Intersystem Data Exchange Operations

Before leaving the subject of DICOM conformance, it is important to repeat a key element of this separation of systems strategy; "the mini-PACS and Enterprise Archive must be able to bi-directionally exchange all types of data objects using DICOM SOP Classes". While many of the next revisions of PACS software will become considerably more DICOM-conformant, there will remain a significant number PACS that will not communicate with their own archive using DICOM operations, much less communicate with a foreign DICOM archive using DICOM operations. There are a number of systems that use a simple FTP command to move study data to their archive server. There are a few systems that do not use DICOM Query/Retrieve to recall priors from their archive. There are a few systems that can optionally use DICOM commands to transfer studies to a foreign DICOM Archive, but cannot use a DICOM command to automatically retrieve those studies from the foreign DICOM Archive.

As previously stated, this is one of the key reasons why the standalone DICOM Archive failed to catch on six years ago.

The concept of an Intelligent DICOM Archive is not new. Emageon, InSiteOne, even General Electric Medical Systems developed standalone archives over six years ago. In terms of market penetration, none of these standalone archive solutions have been very successful.

One reason for this lack of success is simple. Most of the PACS that pre-date the year 2000, and indeed many of that generation that are still around today are not designed to fully interoperate with a foreign archive. They may claim DICOM conformance, but they are lacking in one or more of the four DICOM conformance areas being discussed in this section. In short, many of this older generation of PACS are effectively "proprietary". Only some of the study data can be exchanged with a foreign archive, and the rest is secreted in the PACS. The PACS can send study data to a foreign DICOM archive, but it cannot automatically retrieve it back.

So insisting on 100% DICOM conformance in data objects, headers, compression scheme and bi-directional data exchange between mini-PACS and Enterprise Archive, should solve most of the connectivity problems between the Enterprise Archive and the next generation of mini-PACS software. There are some workarounds including tag morphing

that can overcome the system intercommunication issues between our new Enterprise Archive and the older generation of PACS that are still in operation.

Yet DICOM conformance alone and clever header morphing alone does not preclude the need to migrate the study data every time a mini-PACS is replaced with another.

Granted InSiteOne is providing long-term archiving for millions of images, but each time one of their clients changes PACS, the study data must be migrated from the InSiteOne archive back through the old PACS, into the new PACS, and back to the InSiteOne remote archive. The same is true for all of the other standalone archive products in today's market. While the most common reason for this data migration is the lack of DICOM-conformance in either the old or the new PACS, and therefore the need to extract data objects and morph the study headers, the other major reason is the proprietary nature of almost all PACS Directories.

The PACS Directory Problem

Most vendors in the medical imaging market, whether the product is an imaging modality, a PACS, an Archive, or a Teleradiology system, refuse to disclose their system's Directory Data Dictionary or its Schema. These are the keys to understanding what patients and studies are in the system, how and where study-related information is stored in the Directory, and the file system pointers that describe where the study data is actually stored in the attached storage solutions. So even if system A and system B were completely DICOM-conformant in all four areas described above, the data would still have to be migrated from A to B in order for system B to "learn" all of the patients and studies stored in system A, and in order for system B to create its own file system pointers. So whether the study data is stored on the PACS or on a combination of PACS and foreign Archive, data migration will continue to be a requirement every time you bring in a new system. In this case, the physical migration of the data is the only permissible method for building the new Directory. Obviously the best solution would be for vendor A to open its Directory to vendor B so the Directory and pointer specifics could simply be transferred between the old and the new Directories. The new Directory would then know about all of the studies and know exactly where the data was located on the existing storage media. There would be no need to physically migrate the data. The new system could simply take over the existing storage solution and manage all of the data stored on it. Unfortunately most PACS vendors are reluctant to open their Directories, as if the disclosure of this information to their competitors would lead to their utter and complete ruin. This specific example of propriety in a PACS is perhaps the most damaging and the most important problem to overcome if the organization is ever to take ownership of its data.

There are now three main questions whose answers will constitute the design strategy of the Enterprise Archive.

- 1) "What can be done to eliminate the need to migrate data every time a facility changes PACS?"

- 2) “What is the key to the ‘Never Again’ strategy, where the data is migrated once, and Never Again?”
- 3) How can the proposed Enterprise Archive accommodate those early generation PACS that will remain DICOM-weak in one or more key areas?


The real key to the new Enterprise Archive strategy and the all-important “Never Again” strategy is to once and for all take full ownership of your study data. The only way to accomplish this is what I am going call the “Paradigm Shift”.

The concept of a twenty-year Archive does not work if the Enterprise Archive is placed at the rear-end of the data path. The conventional data path of Modality – PACS -- Archive will simply lead to repeated data migrations of increasingly larger data volumes. The only way to break this cycle is to place the new Enterprise Archive in the middle of the data path. This positions the Enterprise Archive to become the Master Directory and appropriately reduces the importance of the PACS Directory.

The Paradigm Shift

The real key to the “never again” strategy is to shift the roles of the PACS Server and the Archive Server in the system architecture.

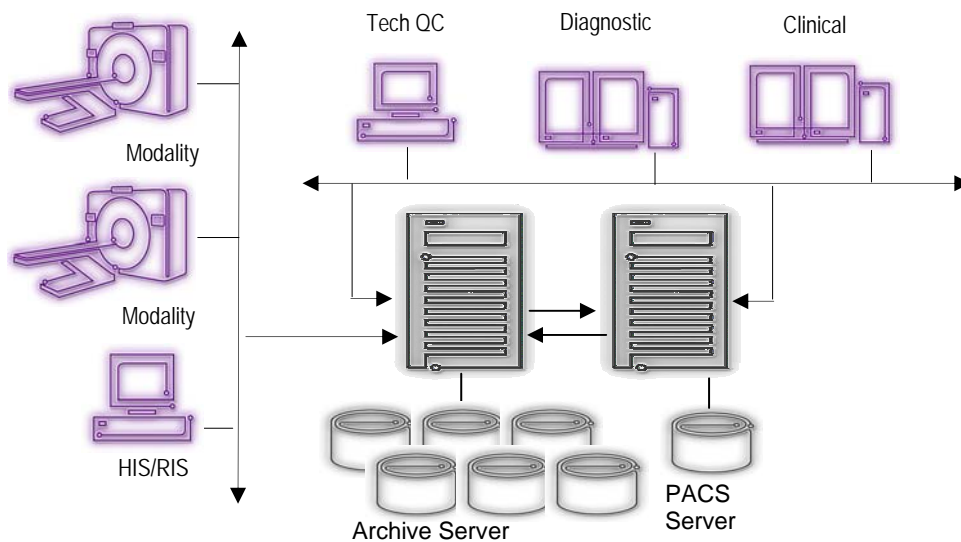
In a conventional PACS, the **PACS Server** is the “centerpiece” of the system architecture.



The PACS Server receives HL-7 feeds from the HIS/RIS. The PACS Server creates DICOM Modality Work List for the modalities and creates work lists for all classes of users. The Modalities forward new study data to the PACS Server, where they are

temporarily stored on the local working cache. The PACS Server builds the Master Directory. The PACS Server supports all classes of Display Station: Diagnostic, Clinical, and Tech QC. The PACS Server retrieves relevant prior studies from its Archive Server. The PACS Server forwards the new studies to its Archive Server once they have been interpreted (status changed from unread to read).

When the **Enterprise Archive** is the “centerpiece” of the system architecture, the Enterprise Archive receives HL-7 feeds from the HIS/RIS and creates DICOM Modality Work List for the modalities and (optionally) work lists for the Technologists. The Enterprise Archive passes the HL-7 feeds on to the PACS Server, which continues to use these feeds to create work lists for the Diagnostic and Clinical users.



The Modalities forward the new study data to the Enterprise Archive, where it is temporarily stored as an unread study. The Enterprise Archive builds the Master Directory. The Enterprise Archive could optionally support the Technologist QC Displays.

The PACS Server’s responsibilities are then reduced to merely supporting the Diagnostic and Clinical Displays.

As soon as the new study has passed the QC process (executed in the Archive), the Enterprise Archive retrieves relevant prior studies stored on its disks and forwards the new study and its relevant priors to the PACS Server. Once the unread study has been interpreted (status changed from unread to read), the PACS Server then returns the new study and any DICOM based study objects created during the interpretation to the Enterprise Archive. The PACS Server builds its own Directory, but this Directory merely supports the Diagnostic and Clinical Display users.

The Enterprise Archive would apply user-defined Information Lifecycle Management rules to the study data as it ages, using DICOM tags to determine when it should be migrated to new or different media and when to purge it from the system altogether. The PACS would continue to store on its local disk cache the new studies and their relevant priors for a reasonable period of time, in order to support access by the physicians. The PACS would manage its on-line study data using its conventional storage management schemes, eventually purging the data.

Changing the data path to Modality – Archive – PACS, places the Enterprise Archive in a position where it can create the Master Directory and manage the study data from its very beginning and throughout its lifecycle (for as long as twenty years). This also places the Enterprise Archive in a position where it can overcome a number of intersystem connectivity issues that are primarily caused by a lack of DICOM and header uniformity among the PACS systems.

It is worth noting that the Archive-centric data flow described above should work with any PACS, since all PACS systems are DICOM-conformant on their front end as a matter of necessity. As for those PACS that are not particularly DICOM-conformant on their rear-end (to the archive), the Enterprise Archive could simply automatically query the PACS every few minutes for any study data whose status was recently changed from unread to read. Most PACS can respond to a DICOM Query/Retrieve command.

As for those PACS that do not archive some of the study data objects with the image data, it is interesting to note that many of these very same systems typically do return all data objects as DICOM objects when they respond to an ad hoc DICOM Query. So it would seem almost any PACS can become more DICOM-conformant, if they are assigned to the proper place in the data flow...at the end.

As I have already suggested, exchanging data between the Enterprise Archive and the PACS is not without its complications.

While there is adequate support in most PACS Servers for the appropriate DICOM SOP Classes that would be used to physically exchange the data with the Enterprise Archive, we have already discussed the problem caused by PACS Servers that do not support the DICOM version of some important meta objects that might be associated with the study, most notably GSPS and KIN. If the Enterprise Archive supports DICOM GSPS and KIN, the technologist would be able to create a Presentation State (Window/Level setting) for the CT images at the CT console and pass this set-up along with the images to the Enterprise Archive. The technologist could also convey some significant clinical information about the study or the patient to the radiologist by creating a brief text note (KIN) at the QC station and storing that with the images on the Enterprise Archive. The Enterprise Archive would then be able to pass these data objects along with the image data to the PACS Server. But if the PACS Server does not support these two DICOM data objects, the radiologist would not be able to access and display the GSPS and KIN created by the technologist using the Diagnostic Display Station tools. Correspondingly, if the radiologist creates and saves a Presentation State, flags key images, and types a text

note using the display station tools, that information might be stored on the PACS Server, but if they are not stored as DICOM objects in the proper header tag, they might not be passed back to the Enterprise Archive, or they might not be recognizable. In the best possible scenario, both the Enterprise Server and the PACS Server support all data objects as DICOM objects and both are able to use the appropriate DICOM SOP Classes to exchange these objects. Of course the best possible scenario will probably not exist for awhile.

Header Tags and the fine art of “Tag Morphing”

A more complex issue is related to non-conformance in the use of Header Tags. For example, if the PACS Server uses a Private Tag as well as a proprietary data format to store the study status (unread/read), how will the PACS Server differentiate the new study from the associated priors that are passed from the Enterprise Archive? Once the study has been interpreted and sent back to the Archive, how will the Archive know that it has been interpreted? This use of Private Tags and proprietary data formats is very prevalent in most installed PACS and it is unrealistic to believe that a wave of conformance will sweep over the industry any time soon. So how is this problem going to be resolved?

Quite simply, the Enterprise Archive has to know enough about the Public and Private Tags used by each PACS vendor to know where key information is stored in the header and how to translate the proprietary format. In our example, the Archive has to be able to locate and recognize the study status as unread or read. Regardless how the Archive vendor chooses to use DICOM within the Archive, the Archive must be able to convert all of the important information associated with an image into the format and Tag location used by a specific PACS Server. The Archive will “morph” the header of the outgoing study data into the header format and Tag location used by the PACS Server. The Archive will also have to locate and translate in the incoming study data (coming back from the PACS) all of the study-related information created in the PACS. The Archive will “morph” the header coming back from the PACS into the format and Tag location used internally in the Archive. It is this clever use of “Tag Morphing” (translation of meta data objects and formats at the level of the header Tags) that will allow a central Enterprise Archive to exchange data back and forth between disparate PACS.

Tag Morphing is that methodology that will allow the Archive to customize the study data being sent to the PACS, so it can be recognized by the PACS. The same methodology allows the Archive to interpret the study data coming back from the PACS and recognize any changes that were made by the radiologist using the Display Station tools. This is how the Archive will recognize the change in study status from unread to read, as well as recognize any new presentation states, key image flags and key image notes. The Archive can then store these changes (in its internal DICOM format) as the “then current” version of the study, and pass them back to the PACS (in its DICOM format) when they are required as priors with the next new study.

If the art of Tag Morphing is taken one step further, it might be possible to locate and morph non-DICOM data objects stored in the image header into their proper DICOM format. This process would make it possible for the Enterprise Archive to actually locate and translate non-DICOM Presentation States and Key Image Notes that are at least stored by the PACS in the image header and convert them into DICOM-conformant formats while they are stored in the Archive.

The art of Tag Morphing was all but non-existent three years ago, and once again this lack of experience contributed to the relative lack of success of standalone DICOM Archives. But today there is actually quite a lot of experience with Tag Morphing that has been gained through the recent surge in data migration between PACS. Vendors have had to learn how to locate and translate information stored in private tags in order to complete study migrations from a legacy PACS they are supplanting.

The QC Option

Changing the data path to Modality – Archive – PACS means that the new study data will first arrive in the Enterprise Archive. The Archive has access to the HL-7 feeds and uses them to create Modality Work List, so naturally the Archive should automatically reconcile this new study against the Order. But how would the exceptions be handled, and how would a user manually reconcile those exceptions? For that matter, what is the best solution for all of the study QC issues that require manual intervention?

Most of the current generation of PACS now include fairly sophisticated Tech QC tools that include verification and editing of study header information, study splits and merges, image deletion, verification and editing of hanging protocols, creation and saving of presentation states, key image flags and key image (text) notes, and document scanning. But many older generation PACS do not have such sophisticated tools, or they are based on expensive dedicated display stations. As long as the Enterprise Archive is going to perform the initial study reconciliation, why not have the option of performing all of the study QC processes using a Tech QC tool kit supported by the Enterprise Archive? Not only would this be the logical solution to how the exceptions are accessed and corrected, but this approach would mean that all study QC, whatever it takes to complete a study and prepare it for interpretation, would be completed in the Enterprise Archive, before it is forwarded to the PACS. This approach would effectively stabilize the Study QC process for multiple versions of PACS.

The Deployment Plan

The following is a step-by-step summary of the Deployment Strategy for an Archive-centric configuration.

- 1) Deploy the Enterprise Archive configured with sufficient storage to manage all existing study data and at least the next year of new study data, assuming the J-PEG 2000 loss-less compression scheme.
- 2) Interface the Enterprise Archive to the incumbent PACS via Gigabit Ethernet.

- 3) Interface the Enterprise Archive to the HIS/RIS in order to access the HL-7 feeds. This might be done by either tapping a second HL-7 feed from the HIS/RIS, or intercepting the HL-7 feed going to the PACS and passing it through the Archive, and on to the PACS. Download all of the orders corresponding to study data currently being managed by the incumbent PACS from the HIS/RIS to the Archive.
- 4) Begin migrating the existing study data from the incumbent PACS to the Enterprise Archive, during off-hours. If possible, extract whatever non-DICOM information is stored in Private Tags in the image header or stored in the Directory database. Convert this non-DICOM information into DICOM formats and store it in the appropriate Tags used by the Archive.
- 5) Reconcile the migrated study data against the original orders retrieved from the HIS/RIS.
- 6) Manually reconcile any exceptions discovered by the Archive's reconciliation software.

Upon completion of the data migration phase...

- 7) Redirect all of the imaging modalities from the incumbent PACS to the Enterprise Archive.
- 8) Activate the Enterprise Archive's DICOM Modality Work List manager.
- 9) Migrate the most recent 3 to 12 months of study data to the most reliable, best-performing storage media already interfaced to the incumbent PACS, and interface any remaining useful storage media interfaced to the PACS to the Enterprise Archive. Retire any storage media interfaced to the PACS that is either unreliable or expensive to maintain.

The New Work Flow

The following is a step-by-step summary of the new work flow in an Archive-centric configuration.

- 1) The Enterprise Archive uses the HL-7 feed from the HIS/RIS to create DICOM Modality Work List for each of the imaging modalities.
- 2) The Enterprise Archive uses the Order information to determine the existence of any relevant priors for each of the new studies ordered.
- 3) Since the Archive knows the amount of storage directly attached to the mini-PACS it knows about the oldest relevant prior that would already be on the mini-PACS. The Enterprise Archive can thus identify which relevant priors will have to be moved from the Archive to the mini-PACS.
- 4) The Enterprise Archive pre-fetches any relevant priors that it knows are not already on the mini-PACS and prepares to move them to the mini-PACS local disk cache.
- 5) The Enterprise Archive performs any Tag Morphing required to convert the header information in the relevant prior images to the header format and Tag location that is recognizable to the mini-PACS.

- 6) The Enterprise Archive performs a DICOM move and transfers the relevant prior images to the mini-PACS local disk. The mini-PACS recognizes the prior images as studies that have already been read.
- 7) The Technologist accesses the Modality Work List created by the Enterprise Archive and forwarded by the Archive to the Modality or its external interface.
- 8) The new study is performed and the new (image data) is forwarded by the imaging Modality to the Enterprise Archive.
- 9) The Archive automatically reconciles the new study against the original order and confirms the accuracy of the Patient and Study information contained in the study header. Exceptions are tagged and posted in a Technologist QC work list.
- 10) The Technologist has the option of manually reconciling the Exceptions [a] using the QC software provided by the Enterprise Archive, or [b] waiting for the study to be forwarded by the Enterprise Archive to the mini-PACS and using the QC software provided by the mini-PACS.
- 11) The Technologist has the option of using the QC software provided by the Enterprise Archive or the QC software provided by the mini-PACS to manually perform any edits to the header information, study splits, mergers, image deletion, etc.
- 12) The Enterprise Archive performs any Tag Morphing required to convert the header information in the new study images to the header format and Tag location that is recognizable to the mini-PACS.
- 13) The Enterprise Archive performs a DICOM move and transfers the new study images to the mini-PACS local disk. The mini-PACS recognizes the new study images as studies that are “unread”.
- 14) The mini-PACS automatically reconciles the new study images against the order, which should be an automatic approval if they were automatically or manually reconciled in the Enterprise Archive.
- 15) The mini-PACS adds the studies that pass QC to the Radiologist Work List.
- 16) The mini-PACS executes its automated relevant prior pre-fetch process and identifies all relevant priors both previously stored on its local disks and those recently transferred there by the Enterprise Archive.
- 17) The Radiologist views the new studies and the relevant priors using the Diagnostic Display software associated with the mini-PACS. The creation and saving of Presentation States and Key Image Notes is optional.
- 18) Following interpretation, the Radiologist changes the study status from unread to read.
- 19) Sometime following the change in study status, the mini-PACS executes its archiving application, which results in using a DICOM SOP Class to forward the new study image data to the Enterprise Archive. Optionally, if the mini-PACS cannot execute a DICOM Move to the archive, or if there is any reason why the “read” study should be pulled from the mini-PACS, the Enterprise Archive would execute an automated DICOM Query/Retrieve for recently read studies and the mini-PACS would respond appropriately.
- 20) The Enterprise Archive performs any Tag Morphing required to convert the header information in the study images forwarded from the mini-PACS to the header format and Tag location that is recognizable to the Enterprise Archive.

- 21) If possible the Enterprise Archive might extract and convert any non-DICOM meta data that was not automatically transferred by the mini-PACS to the Archive. This extracted and converted data would then be stored by the Archive with the study's image data.
- 22) The recently read studies and their associated priors would remain on the mini-PACS Local Disk cache, where they can be accessed by the referring physicians using their respective display station software.
- 23) The mini-PACS regularly purges its Local Disk using whatever disk management application it might support.

Summary

Being Proactive about the inevitable Data Migration from PACS A to PACS B is an easy concept to grasp. Data Migration is inevitable. It is expensive. It is time consuming. It is a problem growing larger by the day, as study volumes, slice counts and image matrices expand. Migrating the data from A to B sooner rather than later makes sense.

One key to the "Never Again" strategy would obviously be to build a DICOM Archive, whose architecture would allow it to evolve through software upgrades and hardware refreshes over twenty years.

The many problems encountered when attempting to integrate a DICOM-conformant Enterprise Archive with a conventional PACS, regardless the level of DICOM-conformance associated with that PACS, can largely be resolved by placing the Enterprise Archive in the middle of the data flow, thus positioning the Enterprise Archive to create the Master Directory, and incorporating in that Archive a library of Tag Morphing routines designed specifically to overcome the major dissimilarities between the systems.

There is a great deal to be gained by this Paradigm Shift. Taking ownership of your data is key among them, but there is also the advantage of a much better way to manage the lifecycle of that data. There is also a very real possibility that your total cost of ownership attributable to managing and distributing medical image data over the next 20 years will be considerably lower than would otherwise be the case. A larger percentage of the costs will be invested in the more stable Enterprise Archive and an increasingly smaller percentage will be assigned to the PACS, that piece that is likely to be replaced every 5 to 7 years. Of course, there are the significant savings to be had by not having to migrate the data ever again. I believe that the separation of Archive from PACS and this specific Paradigm Shift will result in a significant reduction in the overall cost of ownership, because it will eliminate vendor lock, and that nearly always leads to more competitive pricing. This time around the DICOM Archive will be successful.

In short, be proactive. The time to build this new archive is now. Time to migrate study data is now regardless how new the current PACS is, before the migration tasks gets any larger, before you have to add any more storage to a system that will shortly be replaced.

APPENDIX A - Specifications and Requirements

Following are the three key categories of requirements that lay one on top of the other. From bottom up they are:

- State-of-the-Art Storage Solution – the hardware platform and architecture conducive to long-term stability and evolution.
- Intelligent Storage Management Software – the software layer that provides Information Lifecycle Management and flexible administrative tools.
- Intelligent and Automated DICOM Data Management Solution – the DICOM-conformant software application that supports the initial data migration, the HL-7 operations, the pre-fetching of priors, the DICOM Moves and Finds, as well as the automated and interventional QC applications.

Following are specific applications, features, and issues that would determine whether a particular archive or storage solution would be appropriate for your Enterprise Archive-centric configuration. Whether you choose to use a formal Request For Proposal process or not, it would be advisable to request a written response from the vendor to all of these items. Following each question is a comment offered by Gray Consulting (GC).

Section 1 - Enterprise Archive Subsystem and related Issues

Q1. In order to understand the architecture of the proposed Enterprise Archive, please select one of the following descriptions (or combination of descriptions) that best describes your Enterprise Archive Server. If none of these descriptions is a match, please provide your own high-level description of the server architecture.

1. The core applications are combined into a single software package that runs on a single server.
2. The core applications are separated into two or more different software application packages and an individual server is dedicated to each of these different software packages.
3. The core applications are combined into a single software package and a copy is installed on each of two servers that are running in tandem, with so-called “heart-beat” monitoring.
4. The core applications are combined into one or two identical packages and those identical packages are installed on at least two separate server platforms that are configured as a cluster.

GC Comment: The clustered computing concept is the current thing in the computing world. Running multiple copies of a software package on multiple computers arranged in a cluster is cost-effective way to attain both performance and reliability. The individual servers can also be configured with multiple processors and power supplies. The loss of a complete server may have a minor effect on performance, but the “system” is still functional. The cluster concept also provides a simple approach to scaling the system. A

smaller system would require fewer servers, but any size system can easily be expanded to meet future performance requirements by simply adding more servers.

Any Archive system being sold today is likely to feature one of the above Server Subsystem architectures. If properly configured, any of these architectures can provide good performance and reliability. Here are some things to think about when evaluating the architecture of the vendor's Enterprise Archive Server Subsystem:

- The type of system architecture is an indication of the age of the design. The Cluster concept is the newest thing. An Archive with this architecture must therefore be relatively new. That means that there is a good chance that the vendor is not going to re-invent the system in the near future and require you to purchase a lot of new server hardware. Remember, the strategy is to build an archive subsystem that can evolve over many years. The key is to start with the latest technology, and an architecture that is easily expanded to meet growing demands. Growth through the addition of modular servers seems like a strategy that has some legs.
- An Archive configuration featuring numerous Servers, each dedicated to an individual application, such as Image Acquisition, Directory Database, RIS interface, and Archive Services, is clearly an old design. There is every reason to believe that this system is being redesigned and the upgrade to the vendor's next generation will require expensive replacements or reconfigurations of the servers.

Q2. Organizations that are comprised of multiple facilities would benefit from an Enterprise Archive architecture that deploys a Local Archive subsystem in each facility. Does your configuration include a standalone Local Archive Server for each Imaging Facility?

GC Comment: A single, central Enterprise Archive would place enormous pressure on WAN network performance and uptime. A distributed or "federated" architecture would be far more reliable and robust.

Q3. It is a requirement that the various server subsystems be configured with sufficient redundancy to provide 99.9% uptime. Please attach to this proposal the Bill Of Materials for each proposed Server Subsystem and indicate the name of the Server and the uptime guarantee associated with each proposed Server configuration.

GC Comment: Pay careful attention to the BOM for all of the servers in the proposed configuration. Look for as much redundancy in processors, power supplies, network interfaces, interfaces to attached storage, etc. as you can afford.

Q4. What is the Operating System (Vendor, Model, Version) of the various Server subsystems in the proposed configuration?

GC Comment: When it comes to an Enterprise Archive, the issue of Operating System is critical. While many of the newer system designs are embracing Windows in the servers,

there is an argument to be made for using Linux (a version of Unix), if for no other reason than it is not as susceptible to viruses. So the use of Linux in the Enterprise Archive Server subsystem is not an indication of age, as much as it is an indication that the vendor does not yet trust Windows, or does not want the problems associated with using the Windows OS on the servers.

Q5. If any of your server subsystems are configured as a Cluster, please identify the Cluster Technology you are using (Linux, Windows, etc.).

GC Comment: The Cluster technology being used to build a clustered server is worth some consideration as well. This is another Linux versus Windows issue. Clearly there are security and performance issues with Microsoft's Cluster technology that are avoided or better addressed with Linux Cluster technology, and this is yet another reason why some vendors have chosen to use Linux over Windows in their Archive server configurations. The IT department should give this issue the appropriate consideration.

Q6. What Directory Database Management software (Oracle, Sybase, SQL etc.) and Version is featured in each of the various Server subsystems in the proposed configuration?

GC Comment: The Directory database is the log file of the Enterprise Archive. The Directory database is the listing of all patients and their studies, and all data objects related to the studies. For every element in the Directory there is a pointer to the File System that in turn keeps track of where on the Storage Solution the data object is actually stored. The Directory database Schema is the technical description of how the Directory actually organizes this patient, study, series, image content and the respective pointers to the file system. There are a variety of commercially available Directory database products available to the Archive developer including Oracle, Sybase, Microsoft SQL, IBM DB2 to name a few. Oracle has been the most widely used, because it is very feature rich, stable, and scaleable, but it is relatively expensive. MS SQL is becoming more prevalent, especially in those systems that are running Windows OS on the server platforms, because it represents another opportunity to cut costs. I believe the issue is one of balancing Costs versus Performance and Reliability.

Q7. Will the vendor contractually guarantee that the customer will have access to the Data Dictionary and Schema of the Enterprise Archive Directory database, if it is ever necessary to replace the vendor's Enterprise Archive with another vendor's solution, in order to facilitate the population of the new system's Directory with the contents and file system pointers of the vendor's system?

GC Comment: This is a rather critical issue, as failure to obtain access to the Data Dictionary and Schema would make it necessary to migrate all of the data stored on the Enterprise Archive to a replacement system, and one of the major motivations for building the Enterprise Archive is to avoid future data migrations.

Q8. Which of the following best describes your proposed Enterprise Archive system configuration? Please describe any discrepancies or variations.

1. There is only one main Directory database running on one Main Enterprise Archive Server subsystem. This single Directory manages all patient and study information for the Enterprise. If there are any Local Facility Archive Server subsystems, these Local Servers require access to this Main Directory to support local study acquisition and image data management.
2. The Main Enterprise Archive Server subsystem and each of the Local Facility Archive Servers have their own Directory database. Each Local Directory manages all patient and study information for that facility. The Main Directory manages all patient and study information for the Enterprise. Each facility can function autonomously with their own Local Archive Server.
3. All facilities in a multi-facility Health System have their own Local Facility Archive Server subsystem, and each of these servers have their own Directory database, yet each Archive Server knows about the other Servers in the Enterprise. The multiple Directories are effectively combined to form a single virtual Enterprise Directory, so a new study ordered at any one facility will trigger a relevant prior search in all Directories.

GC Comment: Perhaps more important than the choice of Directory database software itself is the Directory Server architecture of the Enterprise Archive configuration. The key to this issue is understanding that each Local Archive subsystem's access to study information is dependent on access to the Directory database and the Directory Server. If a Main Directory database is unavailable, either through a Directory software or Directory server failure, all of the individual Local Facility Archive Servers (without local Directories) would effectively be "down". New studies cannot be acquired and relevant prior studies cannot be identified. A failure of the network connecting a Main Directory Server to a Local Facility Archive Server subsystem will also prevent the acquisition of new studies or identification of relevant priors.

A single Main Directory Server can meet the performance and reliability requirements of a single imaging facility, provided the server hardware itself is configured with sufficient processor power and redundant components. The network reliability can also be addressed with redundant network runs and network switches.

Enterprises that have multiple imaging facilities separated by Wide Area Network connections need to think about the dependency on a single Main Enterprise Archive Server for functionality. Scheduled or unscheduled downtime of the Main Enterprise Archive Server will affect all facilities. Failure of a Wide Area Network connection will affect the connected facility.

An alternative to configuring a fail-proof Enterprise Archive server and doubling up on the WAN connections is configuring the multi-facility system such that each facility has its own Local Facility Archive Server that in turn has its own Local Directory database. In a multi-Directory configuration, each Local Facility Archive Server handles that

facility and all of the facility Directories forward their contents to a Main Directory Server that manages the Master Directory for the Enterprise. All of these individual facility Directories are automatically synchronized with the Master.

Alternatively, in some “federated” architectures, there is no Main Archive Server. Each Facility has its own Local Facility Archive Server and each of these individual Facility Archives work in concert to create a virtual Enterprise Directory.

Regardless the architecture, a Local Facility Archive Server searching for priors related to a new order goes first to its local Directory Server, but the system will automatically search the Master Directory or all of the other federated Archive Directories for relevant priors. Relevant prior search algorithms can be filtered to represent specific facilities, or can be opened up to allow searches of the entire Enterprise.

In order to be truly useful, the architecture of a Multiple Directory Archive must support the exchange of all study data between all the Archive Servers in the configuration.

Q9. Assuming there are multiple Archive Directory Databases (one for each facility), how frequently are each of these individual Directories automatically synchronized with one another during the course of the day?

GC Comment: It is also very important that the various Directory databases are automatically synchronized with each other. In most Multiple Directory architectures, the main or Master Directory server, which maintains the Master Directory of all the patients and studies across the enterprise, is located in the Main Hospital or its Data Center. Each smaller facility has its own Local Directory Server, which manages the Directory of that facility. The Local Directories should constantly be in synch with the Master Directory.

Q10. Please confirm that each of the Archive Server subsystems support all of the following DICOM applications:

1. Store
2. Query/Retrieve
3. Verification Service
4. Performed Procedure
5. Presentation of Grouped Procedure

Q11. It is highly desirable that each Enterprise Archive Server subsystem be able to support and exchange

• Grayscale Softcopy Presentation State (GSPS), Storage SOP Class as a SCU and SCP:
UID 1.2.840.10008.5.1.4.1.1.11.1

• Key Image Notes (KIN) Key Object Selection Document (SR) Storage SOP Class as a SCU and SCP: UID 1.2.840.10008.5.1.4.1.1.88.59.

Do all of the proposed Server subsystems support and exchange the DICOM version of GSPS and KIN?

GC Comment: These are two important DICOM objects. It is imperative that the Enterprise Archive support these Objects and their related SOP Classes.

Q12. The objective of an Enterprise Archive strategy is to store study data from all affiliated imaging services on a single consolidated Main Data Storage Solution based on spinning disk media and Loss-less data compression.

1. Please describe your proposed Main Data Storage Solution(s) for the Main Enterprise Archive Server including Vendor and Model names, usable storage capacity, and approximate compression ratio assumptions.
2. Please indicate all of the SAN or NAS Storage Solutions that you have certified as compatible with your proposed Enterprise Archive solution.
3. Do you offer the option of purchasing the Main Data Storage Solution(s) separately, directly from the OEM or a certified re-seller? What is your proposed cost of interfacing to this Storage Solution(s) if it is purchased separately?

GC Comment: Determining the most appropriate Storage Solution for the Enterprise Archive is a complicated IT issue, because there are so many options such as spinning magnetic disk versus near-line removable media libraries. Is the spinning disk solution directly attached to the Enterprise Archive Server and dedicated to the Archive application, or is the spinning disk solution connected a Storage Area Network (SAN) or a Network Attached Storage (NAS) that is shared among Radiology, Cardiology and perhaps other IT applications? You might configure your Primary Storage Solution one way and your Secondary or Back-up (Disaster Recovery) Solution another way. A large percentage of this study data might have to be managed for as many as 20+ years, so there will inevitably be several types of storage in such a deep archive. You will probably want to configure the Storage Solution for the most recent 12 months of study data one way, the Storage Solution for 60+ months another way, and the storage for data as old as 20 years another way. There are some obvious clinical implications in these issues that are related to both performance and reliability, but the technical issues are largely IT issues.

Question 12 is also testing the vendor's ability and/or willingness to validate and support the connection of their Enterprise Archive Server to multiple network Storage Solutions from such vendors as Compellent, EMC², ExaNet, Hewlett Packard, Hitachi, IBM, Network Appliances, and StorageTek. You don't want to be locked into a single Storage Solution for 20 years.

Q13. The preferred file format for storing image data on each of the Archive Servers in the proposed configuration is DICOM format according to Part 10. This would include the use of DICOM conformant J-PEG 2000 compression schemes. Please describe the

file format you will be using and the data compression options that are available in each of the Servers in your proposed configuration.

GC Comment: Another key to the “Never Again” migration strategy is to start storing the study data in a DICOM standard format from day 1.

Q14. Please confirm that there is some automated methodology for identifying and transferring relevant prior study data from the Main Enterprise Archive Server to any Local Facility Archive Servers, and eventually to the mini-PACS Server.

1. What (event, message, etc.) triggers the transfer of the priors (not already located on the Local Archive Server or mini-PACS)?
2. Is the criteria used to determine relevance to the new study user-programmable?

Q15. It would be preferable if the Enterprise Archive System Manager could log onto the System from any available PC connected to the network and access web-delivered or web-enabled Administrator’s software tools. Can the System Manager application software be run as a thin client from any PC, or does it require a specific pre-configured computer platform?

Q16. How are system errors reported to the System Manager, automatic notification to a pager, cell phone, or email capability?

Q17, It is an imperative that the Enterprise Archive be monitored through a remote monitoring program. Which of the following descriptions best describes that program?

1. A Remote Service Center maintains a 24x7x365 network connection to each of the Archive Server subsystems, and a service technician pro-actively monitors server performance and reacts to maintenance issues as they occur.
2. A Remote Service Center maintains a 24x7x365 network connection to each of the Archive Server subsystems, and the system functions are monitored and captured electronically in a service log, but a service technician only reacts to maintenance issues after they have occurred.
3. A Remote Service Center maintains a 24x7x365 network connection to each of the Archive Server subsystems, and the system functions are monitored and captured electronically in a service log, but the service technician only reacts to a service call placed by the user after a maintenance issue has occurred.
4. A Remote Service Center can activate a network connection to each of the Archive Server subsystems, so the system can be monitored on a scheduled basis, but a maintenance issue is diagnosed only after the user places a service call.

GC Comment: Obviously the Enterprise Archive must be considered a critical high-availability component. It must be protected with a pro-active maintenance program.

Q18. It is a requirement that the proposed Enterprise Archive system include a “Test Server” that meets all of the following:

1. The Test Server includes all of the Enterprise Archive Server application modules represented in the Main Archive Server.
2. The Test Servers allows testing of new software releases independent of the Main Archive Server, and without any effect whatsoever on the functionality or performance of the Main Archive Server during this testing.
3. The Test Server has its own test Directory and Data database, so the test process does not insert test data into the Main Archive Directory database.

GC Comment: Software upgrades will require extensive testing on-site, before they are installed on the production Archive Servers. A Test Server with a complete software application set is imperative.

Q19. It is a requirement that the vendor provide a Disaster Recovery subsystem that meets all of the following. Please describe any discrepancies.

1. Back-up for the Directory database (i.e. Oracle, SQL, etc.)
2. Back-up for the Data database (images and study-related work products)
3. Back-up for Reports

Q 20. Please indicate the Storage Solution that you are proposing for the Disaster Recovery subsystem and any optional Storage Solutions that are available.

GC Comment: Exactly what is the DR solution? Is it remote enough to be a true DR solution? Is the solution as simple, inexpensive, but as cumbersome as a remote shelf library of DLT tapes? Is it as expensive, but as robust as a remote mirror of the Main Archive system? If recovery time is balanced with costs, the best DR solution is a probably a remote Archive Server configured to manage something between 12 and 36 months of study data on-line using economical media, and the balance of older study data on media that is still reliable but much less expensive.

Q21. Indicate one or more of the following statements that best describes the proposed Disaster Recovery Solution

1. Responsibility for data center location and system operation belongs to the customer (self-administered).
2. Replacement of any Server lost in the disaster is the responsibility of the customer.
3. Restoration of the data on the replacement Server is the responsibility of the customer.
4. Responsibility for the data center location and system operation belongs to the vendor (vendor-administered).
5. Replacement of any Server lost in the disaster is the responsibility of the vendor and built into the proposal (i.e. insurance policy).
6. Restoration of some percentage of the study data on the replacement Server and shipment to the customer site within a contractually guaranteed period of time is

the responsibility of the vendor. (The remaining study data would be migrated to the replacement Server later, as time permits.

GC Comment: A DR solution for an Enterprise Archive with the potential for storing 20+ years of data is not going to be inexpensive and easy to manage, yet there is a HIPAA requirement for a DR solution. Look for the appropriate trade-offs of cost versus recovery period and ease of maintenance.

Q22. It is a requirement that the Enterprise Archive support a flexible, user-definable media migration application that (once set-up) would automatically migrate study data from one type of storage media to another within the Archive-managed Storage Solution based on Study Date, Study Type, etc.

1. Does the Archive support a user-accessible media migration GUI?
2. Please describe some of the filters that can be selected by the user to set-up the media migration policy.

Q23. It is a requirement that the Enterprise Archive be capable of selectively purging or selectively retaining study data based on user-definable criteria, such as Patient Name, Study Type, Study Date, Ordering Physician, etc.

1. Does the Archive support a user-accessible data purge/retention GUI?
2. Please describe some of the filters that can be selected by the user to set-up the data purge or retention application migration policies.
3. Once set up, is this purge application automatic?
4. Is it also be possible for the system manager to create and execute custom purging routines?
5. Please describe the nature of the automatic and manual operations.

GC Comment: The Archive is expected to manage study data well beyond the lifetime of most storage media, so the system should be capable of automatically executing media migrations based on user-definable criteria. Eventually study data can be purged from the system, so the Archive must also support automatic and manual purge routines. Furthermore, these purge/retention routines should be user-accessible, because rules are bound to change more frequently than you will want to purchase Professional Services.

Q24. It is our expectation that the Enterprise Archive will optionally support Technologist QC tools. In this case, the Archive Server would create work lists or work folders defined by a variety of user-defined Filters, such as Study Status (incomplete, complete, ID problem, ready for QC, ready for read, etc.).

1. Please confirm that the Archive System can use filters to create work lists for the Technologists QC applications.
2. Please describe some of the filters that can be selected by the user to set-up the work lists.

Q25. It is a requirement that the Enterprise Archive application be able to manipulate both the Public and Private Tags in the image header, such that study data can be translated from whatever format is used by the Archive to whatever format is used by the mini-PACS. This technique is commonly referred to as Tag Morphing, and it is seen to be a key to guaranteeing that the Archive can adapt study data to the peculiarities of a specific mini-PACS, while maintaining a more DICOM standard version of the study data on the long-term archive. One specific example of Tag Morphing is the ability to set the study status tag from unread to read, such that the mini-PACS can differentiate between the new studies and the relevant priors.

1. Please confirm that the Archive application supports Tag Morphing.
2. Please provide a summary list of the types of Tags that you can manipulate.
3. Please provide a list of the PACS that you have successfully interfaced to using this technique.

GC Comment: Tag Morphing is a key to making the Enterprise Archive work with any PACS. Failure to appreciate this requirement or a weak description of the capability is an indication that vendor has little to no experience with this critical tool.

Q 26. How many individual copies of the study data are created and managed by the Enterprise Archive? Where and on what type of media is each copy located?

GC Comment: An Archive that stores two separate copies of the study data will consume twice the storage space, but provide twice the data reliability. This is a cost versus reliability issue.

Q27. It is a requirement that the Enterprise Archive automatically verify the integrity of the study data that it is storing on the system. How frequently does the Archive application verify data integrity: [a] automatically and on a regular schedule throughout the lifetime of the data, [b] routinely but the frequency depends on its age, [c] only when it is first acquired, [d] only when it is accessed.

Q28. How is the data that fails the integrity check replaced? Is this process automatic or does it require manual intervention?

GC Comment: Data Integrity is very important. There are numerous ways for data corruption to occur. The Archive should routinely check for data integrity and there should be a reasonably straightforward solution for restoring any corrupt data. You will probably face cost/benefit considerations related to this issue.

Data Acquisition and Study QC Issues

Q1. Are new studies forwarded by the imaging modalities directly to the Main or Facility Archive Server, or is there some type of an intermediary processing server?

GC Comment: The imaging modalities should forward all new study data directly to the Main Enterprise Archive or Local Facility Archive Servers. This eliminates points of failure and simplifies the acquisition process. Separate data acquisition servers are typically signs of older system architecture. These separate servers are often used to convert incoming DICOM data formats to the vendor's proprietary file format, or compress the data using a proprietary compression scheme. To keep costs down they are usually only single processor platforms and therefore represent single points of failure. Use of such servers is probably a negative.

Q2. It will be very useful if scanned documents such as consent forms, patient history, requisitions, drawings, etc. can be stored as DICOM objects within the study file, similar to the way they are handled by the PACSGear application. Where and in what format are these data objects physically stored in the proposed system?

GC Comment: Treating the scanned document as a DICOM Secondary Capture object and positioning it as a new series in the DICOM study is probably the only way to assure compatibility with the specific mini-PACS.

Q3. It is a requirement that the Enterprise Archive system have the means to automatically check/verify the DICOM Store header in the image data being forwarded by the modalities or their interfaces for the correct patient and exam information and flag suspected studies as "exceptions". Assuming that the Archive software supports a manual intervention application that would be used to "fix" the exceptions, how many screens and how many mouse clicks are required to access the "exceptions" folder and correct one study ID?

GC Comment: First issue is whether such manual QC software exists, and the second is how complicated or simple the Graphical User Interface is to use.

Q4. It is a requirement that the Enterprise Archive system be able to forward studies with suspected ID errors to the mini-PACS, clearly marked as studies with ID errors, so the radiologist has the option of viewing or interpreting without delay. What assurances can you provide that your Archive will be able to communicate this study status (exception) to any mini-PACS it might be interfaced to?

GC Comment: If Tag Morphing is the answer, what evidence is offered that the vendor has this experience?

Q5. After correction of a failure in the Modality Work List (MWL) or RIS interface, it is a requirement that the proposed system automatically check the ID of all the studies acquired during downtime and mark any suspect studies.

Q6. It is a requirement that the proposed Enterprise Archive system allow the radiologists to access and read "John Doe" trauma cases before the ID is corrected. Explain how this requirement is met by the Archive.

Once the patient name is known, which systems must be manually updated with the change (HIS, RIS, Archive, mini-PACS)? Which systems are automatically updated once the name change has been made in one system?

GC Comment: Ideally the Name change is made in the RIS and that change is propagated automatically to the Archive Directory(ies) and the mini-PACS Directory(ies).

Q7. If the Enterprise Archive system configuration includes more than one Directory Database (Oracle, SQL, etc.), you should expect that the system would automatically synchronize or update each of the individual Directory databases once an ID has been corrected in one Directory. Does the Enterprise Archive system meet this expectation?

Q8. Does the proposed Enterprise Archive system convert incoming DICOM-formatted data transferred from the imaging devices to a proprietary (non-DICOM) format for storage in the Archive?

GC Comment: The answer should be no. There is no longer any significant advantage to be gained by the Enterprise Archive converting DICOM data forwarded by the imaging devices to proprietary data formats. In fact, any advantage to storing and using proprietary data formats in the Archive is offset by the time and cost related to converting that data back to a DICOM format when it is time to forward the data to the mini-PACS.

Q9. Please provide a list of the specific models and versions of HIS, RIS, Interface Engine, and Master Patient Index (systems) that you have successfully interfaced to the Enterprise Archive.

GC Comment: Placing the Archive in the middle of the data flow, specifically directly behind the imaging modalities means the Archive is going to have to interface to one or more of the subsystems mentioned. It would be useful and reassuring to know the vendor's level of interface experience.

Q10. Can the Enterprise Archive pass the HL-7 feeds from the HIS/RIS or Interface Engine on to the mini-PACS, or will two, separate sets of HL-7 feeds be required to interface to the Archive and the mini-PACS?

GC Comment: The ability to pass the HL-7 feed through to the mini-PACS would save the costs associated with the purchase and maintenance of a second set of HL-7 feeds.

Q11. It is a requirement that the proposed system be able to manage study data created at all of the radiology department(s) (imaging facilities) on the same Main Enterprise Archive Server, even if one or more of these imaging departments is running a different RIS. Furthermore, it is a requirement that all of a patient's studies, regardless where they were performed, and regardless the uniqueness of the Medical Record Number, be organized in a single electronic "folder". Please describe what feature of the Enterprise Archive or what (if any) additional software or subsystem will enable the Enterprise Archive to support this "co-mingling" of study data.

GC Comment: If there are multiple RIS or multiple instances of the same RIS in the Enterprise (meaning that there is not a single unique Medical Record Numbering system across the entire enterprise), it may be necessary for the organization to purchase and maintain a Master Patient Index to combine the same patient's studies from different facilities into the same folder. The Enterprise Archive should have optional software that provides this MPI-like functionality at a fraction of the cost of an MPI. The Archive could then manage all of a patient's studies and recognize that they belong to the same patient. Does the Archive offer a way to do this short of using a MPI?

Q12. Please verify that the Enterprise Archive will be able to receive and process the HL-7 feeds (ADT, ORU, ORM) in order to create and provide DICOM Modality Work List for all of the imaging devices.

GC Comment: A key component of the Enterprise Archive System is the RIS Gateway or Interface. This Gateway or Interface provides a necessary link between the Enterprise Archive (image) system and the RIS/HIS (information) system. This RIS Gateway would support the creation of Modality Work List by the Archive, as well as a number of other Workflow applications. The newer system designs have eliminated the external Gateway box and internalized the software application. The term "broker-less" simply indicates a newer design that has eliminated an external single point of failure.

Q13. Please verify that the Enterprise Archive will be able to acquire and manage the final (signed) radiology report. Please indicate the file format of this report. Can the Enterprise Archive translate the report back and forth between a Text format and the DICOM Structured Report format?

GC Comment: The radiology reports are typically "archived" in the HIS or RIS, but the Enterprise Archive is supposed to survive even these systems, so why not store a copy of the reports in the Enterprise Archive as well? The most open data format would be DICOM Structured Report, but the Archive will probably have to have the capability to translate this report from this format to a text-based format compatible with the mini-PACS.

Q14. One or more of the imaging facilities may from time to time receive studies transmitted electronically or on a CD/DVD from an "outside" facility. It is a requirement that the proposed Enterprise Archive be able to "import" these outside studies and assign them a unique identifier (such as a facility prefix to a study ID).

1. Furthermore, it is a requirement that the proposed Enterprise Archive provide the option of either deleting an "outside" study or archiving the "outside" study based on predetermined criteria applied to the unique identifier created by the Archive.

GC Comment: The outside study should be imported into the Archive, if there is any chance that you will want to add it to the long-term archive. This process would most easily be conducted at a Tech QC station.

Q15. Certain classes of “outside” studies (i.e. based on Facility ID) are not needed after their interpretation. The Enterprise Archive should support either manual or automatic purging of the unwanted studies from the system, once the report has been signed. In this case purging means erasing the image data from the Data database and the study ID from the Directory database. Please confirm that the Enterprise Archive can support this Directory and Data database purging based on user-defined criteria and Facility ID, AE Title, etc.

Q16. Which (if any) of the following types of Cardiology data can be transferred to the Enterprise Archive: Cath Lab, Echo, ECG, Cardiac CT, Cardiac MRI?

GC Comment: The Enterprise Archive concept assumes that each imaging department will share the Archive. Cardiology data objects are different, but many have been assigned DICOM object status. All of the above should be supported.

Technologist QC Option

It is the opinion of Gray Consulting that the Radiology Technologist should play a major role in Study QC. While the System Administrator is responsible for overall system functionality, troubleshooting and maintenance issues, the Technologist is responsible for assembling the study data and all study-related clinical and business information that the radiologist requires to interpret a study. If the Tech uses the RIS terminal to open and close a study, I think the Tech will use the Tech QC Station to perform study QC. I’m putting this issue in this context, because I believe that you will probably need nearly as many Tech QC Stations as you have RIS stations scattered around the department.

Think about all the questions in this section and I believe you will understand why you will probably place a Tech QC station in each of the imaging sections. You need just enough of them to prevent a line of technologists waiting to get their work done.

If the original study is going to be forwarded to the Enterprise Archive first and the Archive is going to reconcile the study header against the order, it seems logical for the Archive to support some level of Technologist QC application software as well. The following questions are designed to see how many QC tools are supported by the Archive. Ideally all of them will be supported, so the Technologist can properly QC the study before it gets sent to the mini-PACS and on to the radiologist.

Q1. Does the Enterprise Archive provide an optional Tech QC application software package that can at the very least be used to edit image headers flagged as “exceptions” by the Archive’s study reconciliation routine?

GC Comment: The technologists are probably better equipped than the System Administrator to quickly resolve any exceptions flagged by the Archive’s reconciliation process. The tools to support this manual reconciliation and data entry should be wrapped in an easy-to-use Graphical User Interface and be part of a QC software

application that can be run on the same basic PC that is already supporting the RIS application in each imaging section.

Q2. Is this Tech QC software a web-delivered or web-enabled application, whereby the master copy is hosted by the Archive Server and the client copy is automatically downloaded from server to client at initial set-up and run thereafter on the local client platform?

GC Comment: It would be cumbersome to have to manually load the Archive QC software on the RIS PCs. The current generation of PACS are based on web-delivered display software. This approach is simple to set-up and maintain.

Q3. Does the optional Tech QC application software package include a document scanning package?

GC Comment: It is advisable to have a document scanning strategy in your imaging department. Since the documents related to the study are often stored with the images, it seems appropriate to scan the documents into the Enterprise Archive and store them with the image data.

Removing paper is a complicated workflow issue. Some of the information on the paper forms currently in use can be captured by the RIS or by data fields in the PACS, but a lot of the information currently captured or entered on paper will probably remain on paper. The simple solution is to continue to use paper in the front end of the study process, but figure out where in the process you can scan the paper. One logical solution that works for a lot of radiology departments is to have the clerical staff scan some of the paper at the front desk or in the file room. Then have the Technologist scan whatever forms they have added information to in the course of performing the study. In short, scan a document at that point in the study process when no additional information is going to be added to it, and have the person who is holding the paper at that point perform the scanning process. Using this strategy, it makes sense to add the scanning software and hardware to each Tech QC Station and to each RIS Station at the front desk and in the film file room.

Q4. It should be possible to use the existing PC's running the RIS software as the host platform for the Technologist QC Station and Document Scanning software. Please state if this is possible, or any reasons why the Tech QC Station requires its own dedicated platform.

Q5. Which of the following major functions are supported by the Technologist QC Station software. Please describe any discrepancies.

1. Patient and Study ID Edits
2. Clinical Info Input; verifying existence, completeness and accuracy of clinical info such as ordering physician, history, reason for study, etc.
3. Image Display

4. Document Scanning
5. Study/Image Deletion; deletion of some or all images from a study
6. Study Merge; merging of two series of images belonging to the same study
7. Study Splits; splitting a single large study into multiple studies with overlapping images
8. DICOM Print SCU
9. Postscript Print
10. Copy of images or studies to CD/DVD burner application.

GC Comment: A complete QC package would support all of these applications.

Q6. Please identify the document scanning application software that can be interfaced to or integrated with the Technologist QC software., and confirm that the scanned document object is in fact a DICOM Object.

GC Comment: The commercial package available from PACS Gear is a solid choice, used by most PACS vendors, but an in-house application is acceptable. Just make sure that the scanned document object is a DICOM Object, preferably merged as a new series into the study. That assures that it will be accessible to the PACS.

Q7. Which of the following work lists or folders can be created for the Technologist Display Station:

1. Exceptions - Studies with detected ID/information errors
2. Incomplete – Suspended studies that await additional images
3. “John Doe” – Studies for unidentified patients
4. Complete – Completed studies that await Technologist QC
5. Device Specific – Work List (Folder) filtered to a specific imaging device (AE Title)
6. Modality Specific - Work List (Folder) filtered to a specific modality (CT)

GC Comment: All of these folders would be useful to the technologist work flow.

Q8. Are the Work Lists, Folders, and all other user preferences supported by the Technologist QC Station user-specific, such that the QC Station in the CT area could be customized for the CT Tech?

Q9. Can the Tech QC Station be used to rearrange the images or series within a study so that the study will work with an existing Hanging Protocol designed for the specific mini-PACS to which the study will be forwarded? Can the Tech then save the rearranged study and does the Archive then substitute the rearranged study for the original before sending it to the mini-PACS for interpretation?

GC Comment: It is unrealistic to expect the Enterprise Archive to create or modify hanging protocols for any given PACS, as hanging protocols are not subject to any standard and are thus quite unique to a given PACS. Nevertheless a technologist would

probably know how to arrange series within a study to insure that they are in the order expected by the PACS hanging protocol, so the Tech QC station should support this feature.

Q10. We would like every Tech QC Station to include the following CD/DVD Copy applications.

1. The option to copy Images and associated Presentation States, Key Image flags and Key Image Notes to the CD/DVD in the DICOM file format
2. The option to copy Images and associated Presentation States, Key Image flags and Key Image Notes to the CD/DVD in a non-DICOM file format
3. The option to copy an image viewer (software) to the CD/DVD
4. The option to copy multiple studies of the same patient to the same CD/DVD
5. The option to copy multiple studies from different patients to the same CD/DVD

GC Comment: The PACS could also be used to support the CD/DVD applications, but it is logical to place these tools in the Tech QC Station associated with the Archive, since this is the tool that the technologist would be working with most of the time.

Q11. Please specify whether the Tech QC software license is an unlimited site license or is based on the actual number of QC stations.

GC Comment: The best pricing option is an unlimited site license fee.



Michael J. Gray

Mr. Michael Gray is the principal of Gray Consulting, a consulting practice established in 1991 to develop a number of consulting services designed to assist the Integrated Delivery Network, individual Hospital, and independent Imaging Center design, choose and deploy the “right” PACS. Gray Consulting has provided these PACS-related consulting services to over 45 Health Systems.

Mr. Gray’s areas of expertise are Market Analysis, Technology Analysis, Strategic Planning, Equipment Utilization, Needs Assessment, Workflow Analysis and Re-design, Business Case Modeling, Vendor Analysis/Selection, and Technical Contract Negotiation. Mr. Gray’s unique perspective has proven invaluable to Radiology and Information System administrators planning the introduction of new Image Management technologies into their imaging departments and hospitals. Mr. Gray routinely speaks to both national and local Health Care Organizations on subjects such as post-PACS workflow design, business case modeling, system deployment strategies, expansion or replacement of data Storage Solutions, and development of Data Migration strategies from old to new PACS.

“ The foundation of my consulting services is an unbiased and completely ethical approach. It assures that my clients receive the full benefit of my experience and expertise without compromise. You should also know that I personally perform all projects...there is no cadre of “associates”.

Mr. Gray has a BS in Biology and Chemistry from Washington University, St. Louis; holds three US patents; and has an extensive bibliography in medical image display and electronic information management systems.

Mr. Gray and his family reside in Novato, California.

Contact Information:

Michael J. Gray

Gray Consulting

P.O. Box 972

Novato, CA 94948-0972

Office: 415-892-0943

Cell: 415-328-2842

Fax: 415-897-0384

Email: graycons@well.com

Web: <http://www.graycons.com>