

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

Separating PACS Servers from VNA ...and then connecting them

By Don Dennison
May 2014

Sponsored by: Hitachi Data Systems and Brocade Communication Systems

Contents

Executive Summary 3

First, what is a VNA? 4

What are some Requirements of a VNA? 4

 Interface Quality 4

 Scalability 5

 Resiliency 5

 Data Integrity 5

 Security..... 6

 Data Access Control..... 6

 Cross-indexing of Patient Records 6

 EMR Integration Capabilities 6

 Sharing Records Across Enterprises..... 7

 Metadata Coercion 7

 Storage Flexibility 7

 Intelligent Routing..... 7

 Information Lifecycle Management..... 8

 Support for Enterprise Imaging..... 8

 Breadth of Data Types Supported 8

Why do PACS servers often fail to perform well as VNA? 8

So, where do PACS servers come up short as VNA systems? 9

 We Have Our Orders..... 9

 Automated Processes for Reading Workflow 9

 A Failure to Communicate 9

 Patient Identity Management across Domains 10

 Limited Storage Technology Options..... 10

 Limited Ability to Manage a Record's Lifecycle..... 10

 Limited Metadata Available..... 10

Deployment Flexibility..... 11

Integrating a VNA with PACS—What to Consider 11

Infrastructure Proficiency 11

External Archive Capability in PACS 11

Prefetching and Relevancy..... 11

Automated Prefetching 12

Manual Retrievals..... 12

Interoperability of Entire Imaging Record..... 12

Additional Notes 13

Content Monitoring 13

Data Synchronization..... 14

The Future 14

Additional Reading 15

About the Author 16

Executive Summary

With the explosion of data occurring in healthcare, many organizations are looking for ways to consolidate costs and improve efficiencies. One of the ways available to providers is the acquisition of a Vendor Neutral Archive. In traditional circles, the VNA represents a way to consolidate medical images from various departments, such as radiology and cardiology. But with the increase in unstructured data from a variety of sources outside medical imaging, a VNA can mean so much more.

Hitachi Data Systems has built its business around archiving. Key to our strategy is managing the content in addition to storing data. By managing the content, the traditional VNA becomes more than just an archive; it becomes the foundation upon which a facility can build clinical analytics as they drive towards population health management.

Hitachi Data Systems and Brocade understand what it takes to deliver a critical infrastructure that can be used throughout the lifecycle of patient data. Following an open standards philosophy, believing that "green" is our future and that reliability can't be underestimated when it comes to patient safety, HDS and Brocade have combined to deliver a best-of-breed Vendor Neutral Archive. Hitachi Clinical Repository meets the needs of healthcare providers now and into the future. Enjoy this whitepaper with our compliments and learn more about what Hitachi and Brocade can do for your organization today.

First, what is a VNA?

Vendor Neutral Archive or VNA is a term used to describe a system that acts as a long term archive and data sharing platform, typically providing image and information storage and management services to one or more PACS (Picture Archive and Communication System) and/or imaging modalities. It may also be used for storage and management of clinical or so called enterprise images. Their popularity has risen out of the desire for a more open and flexible archive system than was typically available as part of the PACS. It is also a common IT component in an enterprise, and even cross-enterprise, consolidation strategy.

In the simplest of deployments, the system provides better archive and data lifecycle management than what a hospital's PACS often does. It may even offer better integration interfaces, with more support for IHE (Integrating the Healthcare Enterprise¹) integration profiles than the PACS. In these environments, the VNA fills functional gaps in the PACS, and—once the legacy data is migrated into it—makes replacing the PACS easier.

In more sophisticated deployments, the VNA often serves multiple integrated systems, acting as a sharing waypoint, likely in addition to acting as the long term archive. It will often provide services to normalize the patient's image record data, which may include cross indexing of patient identities (e.g. using the IHE Patient Identity Cross-referencing, or PIX, integration profile).

By integrating an Enterprise Viewer, it can provide a common yet controlled access to images for users across the enterprise; even directly from within the Electronic Medical Record (EMR) or Health Information Exchange (HIE) system.

VNA have been around for more than a decade, but have risen in popularity due to consolidation among providers, the desire to manage images at the enterprise (rather than buying more and more archive storage in various departmental systems, and maintaining storage management expertise in the departments) in order to reduce costs, and the rapid adoption of EMR and HIE systems (when the operators wish to include imaging records) in markets like the U.S.

What are some Requirements of a VNA?

As with any product, customer priorities will vary, but there are some common requirements cited for a VNA.

Interface Quality

The flexibility and breadth of the data storage, query and retrieval interfaces, such as those using the [DICOM](#) (Digital Imaging and Communications in Medicine) standard, are critical, as the VNA is being counted on to be able to interact with a wide variety of often less flexible systems (e.g. a PACS). It is often used to compensate for the inferior interfaces of a PACS, when integrating with another system. Support for DICOM and [HL7](#) (Health Level 7) standards are essential, but other protocols, such as those defined in IHE XDS-I (Cross-enterprise Document Sharing for Imaging²) and IHE [ATNA](#) (Audit Trail and Node Authentication) integration profiles may also be important, depending on the needs of the buyer.

¹ <http://ihe.net/>

² Also, refer to the IHE Cross Community Access for Imaging (XCA-I) integration profile, which defines how to connect multiple XDS-I communities for record discovery and access.

The term VNA is most often used in the U.S., with some usage in Canada and the UK. Other parts of the world may use a different term for a similar system.

Ideally, the VNA is flexible enough to proxy one transaction protocol to another; an example of this is providing a method for a standard DICOM query (C-FIND) to be translated to the necessary XDS-I query, providing results back as a DICOM query response. Capabilities like these allow systems that support only one protocol/API (e.g. DICOM) to participate in an integrated multi-enterprise environment, without having to support the new protocols/APIs.

Scalability

As more systems are connected to the VNA and more data is stored and accessed, the system needs to be able to scale out significantly to provide desired performance under the substantially increased load. For example, a system may handle storage of 300,000 exams per year well, but may fail to handle 3,000,000 exams per year. Such 10x (or more) increase in expected load often requires more than incremental processing capacity (e.g. adding more application servers), but a completely different architecture. Support for virtual servers makes scaling out processing capacity more cost effective and easier to manage. Support for the use of parallel DICOM associations is desirable to provide data transfer efficiency.

Transitioning to a VNA environment promises many benefits, however, new challenges can arise, such as the need to maintain consistency of policy and performance, regardless of where a workload component resides or moves to. This means ensuring that network, compute and storage services are synchronized throughout the life of the workload. An open, standards-based architecture will help drive seamless integration across all aspects of the infrastructure, including the application layer.

Resiliency

As a critical component in health record management, protection and availability of patients' imaging records must be assured. Availability of affordable Disaster Recovery (DR) and Business Continuity (BC) solutions are an important consideration. Infrastructure now plays an even more critical role in delivering an always available environment, where five 9's availability is simply not enough.

Traditional data center infrastructures were never designed for today's astronomical growth in healthcare interoperability, bandwidth and data access control. Software-Defined Networking (SDN) is a powerful new network paradigm designed to address these exact issues and SDN-ready infrastructures should be considered for all new deployments.

The infrastructure must be available at all times, regardless of circumstance. This level of availability can be achieved through active/active architecture configurations. Additionally, proactive system monitoring, with effective notifications and reporting, are key to measuring and ensuring operational resiliency.

Data Integrity

In addition to system-level protections, more granular record-level protections should be provided to detect and identify undesirable record modification or corruption (e.g. by a defective network, server or storage component). Industry-accepted integrity checks of database data and stored files should be an inherent part of a VNA design.

Security

A VNA should be capable of storing files—so called Data at Rest (DAR)—in an encrypted format; this may require an integrated method to manage the encryption keys³. Most modern database management systems provide options for encrypting the stored metadata. Encryption of data in transit is also desirable. Transport over HTTPS is well understood by IT and normally easily achievable. Secure DICOM communication, referred to as DICOM TLS (a standard), is possible, but not widely supported by common DICOM systems. HL7 communication can also be done over TLS. Auditing of all significant events should be provided; typically, IHE ATNA is desirable (at minimum)

Data Access Control

As a VNA is often used to consolidate content from multiple systems in multiple organizations, a capability to control access is needed. As business agreements and policies are established (and change, over time), the system needs to be able to enforce these.

Cross-indexing of Patient Records

In order to support the sharing of records across enterprises, the VNA must provide some method of indexing records for the same patient even when different patient identifiers (e.g. Patient ID or Medical Record Number) are used. This often involves a combination of query and retrieve interface features, and integration with an MPI (Master Patient Index) system and/or mapping of metadata.

EMR Integration Capabilities

One of the roles of the VNA is often to provide a consolidated integration point to incorporate images into an EMR or HIE system. This normally involves the use of a so called Enterprise Image Viewer integrated with the VNA and into the EMR/HIE. A couple of common approaches are used:

Notification – The VNA sends a message (usually HL7⁴) to the EMR/HIE that a study has arrived and is available for viewing. The message contains key metadata needed by the EMR to launch the Enterprise Viewer. The Enterprise Viewer is launched by the EMR (normally by issuing a URL to the Enterprise Viewer's server).

Query/Retrieve – The EMR will query the VNA, on demand, to discover what data is available for a given patient or patients. If the Enterprise Viewer is logically a part of the EMR, the VNA may also need to be able to return DICOM objects, upon request. Existing DICOM standards such as, WADO⁵, as well as new ones such as WADO-RS⁶ and QIDO-RS⁷ provide standard API options over proprietary methods.

³ This capability may be provided by the VNA application software, or storage technology layer.

⁴ Another protocol to consider is DICOM and its method called Instance Availability Notification; though, as most EMR/HIE do not support DICOM, this may be more useful when integrating with other imaging systems.

⁵ DICOM Supplement 148; Web Access to DICOM Persistent Objects

⁶ DICOM Supplement 161; Web Access to DICOM Persistent Objects by RESTful Services

⁷ DICOM Supplement 166; Query based on ID for DICOM Objects

Sharing Records Across Enterprises

Multiple enterprises or communities may be connected via DICOM/HL7 or XDS-I/XCA-I; however, this can quickly become unmanageable if every system needs to be connected to every other system. As a result, a VNA is a natural “broker” to bridge the local PACS to the outside world. For example, the VNA can act as a DICOM query/retrieve proxy to one or more external PACS upon receiving a request from a local PACS. In this case, the local PACS may not even have direct connectivity to the external PACS. In a more complex configuration such as XDS-I and XCA-I, the VNA can act as a XDS gateway for the local PACS to the XDS Affinity Domain or other XDS communities.

Metadata Coercion

Also referred to as “tag morphing”, this involves the capability of detecting patterns in the data stored (or based on the source), and applying a conditional mapping or modification to the data upon receipt (inbound). The same capability is needed on retrieval of data by external systems (outbound). This may be used to inject some desired text into the record for future use, or to correct some invalid data before archiving. In general, there are a couple of levels of metadata coercion: 1) coercion based on mappings; 2) coercion based on dynamic queries to an external system. An example of the latter is a service that uses metadata from an inbound study to look up patient, or order/procedure info, and apply it to the header at time of store. The typical metadata coercion can deal with less unique values, like those for DICOM attributes Modality or Institution, but could not do Patient ID (for example) mappings other than prefixing/suffixing, or moving a value from one DICOM attribute to another (e.g. Insurance number to be appended as a value to the Other Patient ID Sequence attribute). Some common DICOM attributes where coercion is often required include Body Part Examined, Anatomic Region Sequence and Institution Name.

Storage Flexibility

As a data management and archive system, the interaction with storage technology is one of the primary functions of a VNA system. The flexibility as to how the VNA makes use of different storage technologies is a key value of the system. Being able to choose a valid product from a preferred storage vendor, or being allowed to connect the VNA to an existing storage system, are often desired. Whether integrating or migrating data, evaluate the storage environment to ensure low latency, highest port density and simplified management is in place to maximize overall uptime and performance. Being able to add new storage to an existing system, and move existing data over to that storage, is also important. Being able to make use of specific capabilities or APIs of the storage product, to realize the benefits of these capabilities, is often desirable.

Intelligent Routing

As a system that brokers data among other systems, the VNA often needs to have logic to move record data to an external system; for example, an automatic route (DICOM C-STORE) of desired data to one or more systems. Capabilities to seek include the type of content or source of the content to be used to trigger the route. Also, investigate the scope of the data to be routed—where some systems can route a complete study, in some cases, it is desirable to route only selected series or objects.

New architectures and new protocols are needed to ensure the secure, free flow of information within the various systems. SDN-ready infrastructures will enable organizations to tailor traffic in a new way, allowing the network to push changes based on programmatic control.

Information Lifecycle Management

As a record management system, a VNA needs the ability to be able to manage the lifecycle of the records. This may include managing the availability and location⁸ of a record, based on events or conditions, up to permanently purging records, when policies dictate. The rules expressing these policies should be flexible, and the configuration interface for these rules clear (and access controlled).

Support for Enterprise Imaging

Some PACS include a design limitation that prevents the system from storing patient-level information (e.g. from an HL7 ADT message) unless at least one imaging object is stored; this is because the PACS is optimized for acquisition and reading workflows. If using the VNA to store and manage enterprise imaging data—such as clinical images or diagnostic images from outside traditional departments like Radiology, Cardiology, and Nuclear Medicine—there is often a need to have the VNA store patient information prior to any images being captured. Also, while virtually all PACS, and most modalities (via DICOM Modality Worklists), support some method of using orders (HL7 ORM messages) to ensure valid and consistent (with the Hospital Information System, or HIS) record metadata, enterprise images are often not “ordered”, so requiring order information to validate the imaging data from some sources is not desirable or feasible.

Deciding where to store documents is an important one. While some VNA products offer the ability to store documents (associated with a patient) other options, such as an ECM or EMR should also be evaluated. Considerations include how and where the documents are captured, how they will be accessed, and what data, such as a visit or episode identifier, needs to be associated with them.

Breadth of Data Types Supported

Obviously the more types of data that the VNA can accept and manage, the more useful it is. When evaluating a VNA, careful attention should be paid to the DICOM Conformance Statement; in particular, to the supported SOP Classes and Transfer Syntaxes. Where an EMR or ECM (Enterprise Content Management) system does not meet non-DICOM document management needs, the VNA may be asked to provide management services for this type of data, as well.

Why do PACS servers often fail to perform well as VNA?

The ‘A’ in PACS stands for ‘Archive’; so why don’t PACS meet buyers’ needs for archiving? Truthfully, sometimes they do. If the site has no plan to replace their PACS, has no organizational mergers or consolidations on the horizon, have been able to integrate it with other systems when they need to, and no strategy for capturing and managing enterprise images, a stable, well-managed PACS will often do the trick.

But the above situation is increasingly uncommon. The reality is that the “archive” inside the PACS is just one component, and it does not get the same attention as the reading workflow management, image display and report creation functionality—or at least it shares its attention with these more visible components of a PACS. For VNA vendors, it is their complete focus; and what gets attention, gets better.

⁸ A location means the storage system where the data files are persisted; this method can be used to move data among different storage technologies with varying performance characteristics and cost, as desired.

PACS vendors will often make the mistake of assuming that their PACS server can be a VNA, typically by looking at the trivial similarities such as storing DICOM objects, query/retrieve capabilities, maturity of technology, and maturity in the deployment. But is it not that simple.

The challenge of creating a single application that can serve as both a PACS server and a VNA is not a technical feasibility one, but one of competing design priorities. Developing a product that both optimizes services for managing acquisition and image delivery to PACS clients, as well as manages interoperable imaging records across enterprises, is difficult. But, it can be done if the right amount of resources are invested.

So, where do PACS servers come up short as VNA systems?

We Have Our Orders

First, PACS have been explicitly designed and optimized for the acquisition and reading workflows of the Radiology department. This is good for departmental operations; but, like any optimized design, there are trade-offs. One of these is a common requirement that the PACS receive an order before a received study is accessible in the PACS. This feature is in place to ensure that received image metadata is updated with current patient and order information before being used for diagnostic interpretation. Not all received data will have orders (for example, enterprise images, migrated data, and legacy data shared from systems external to the enterprise). Some PACS include workarounds to remove this requirement per storing system, but not all do.

Automated Processes for Reading Workflow

In a departmental workflow, it is reasonably common for a PACS to flag a study when images arrive after it has been dictated. This exception workflow is intended to ensure that newly arrived images are reviewed to give the Radiologist the opportunity to confirm or revise their findings, based on what the images show. However, a VNA may encounter a similar situation for entirely different reasons. For example, due to queuing behavior within the PACS or a network delay, a portion of a study may arrive at the VNA much later than the rest of the study. Another example is the use of DICOM to encapsulate other documents such as a scanned requisition or diagnostic reports. These DICOM objects may be created by the VNA or sent to the VNA after the study has been dictated. As shown in these examples, a valid and desirable workflow in a PACS may require adjustments in the domain of a VNA.

A Failure to Communicate

Another common shortcoming is in the breadth of data types accepted and managed. Common Radiology and Nuclear Medicine data types, as well as most Cardiology ones, are often supported, but many DICOM SOP classes and Transfer Syntaxes used by other diagnostic and clinical departments are not as common.

It can also be the case that the DICOM interface of the PACS server simply isn't as robust as one would want. This may be evident in the limit of concurrent DICOM associations allowed per PACS server, or the throughput of objects (stored in and retrieved out) simply is not fast enough. Another limitation may be the breadth of query criteria and return keys allowed in the DICOM query (C-FIND); though, PACS that support IHE SWF (Scheduled Workflow), and more specifically the Query Images [RAD-14] transaction it specifies, should provide a reasonable query capability.

Patient Identity Management across Domains

The ability to manage imaging records with patient identifiers from different domains varies wildly among systems from different vendors. Some are capable of managing only a single patient ID (aka MRN⁹); and some of these do not support identification of patients through the paired values of the DICOM attributes Issuer of Patient ID (the name of the system that generated the Patient ID) and Patient ID (the actual patient identifier value), resulting in a risk that two different people with the same Patient ID value (e.g. 123) assigned by two different Issuers (e.g. HIS at hospital A and HIS at hospital B) get treated as one patient in the system.

Some systems support the use of an MPI value in addition to the original patient identifier, allowing them to manage patients from across more than one domain. The capability to either map the MPI value from received data (attribute in the DICOM header or value in the order) or to query an external system (an MPI), as needed, will also vary.

As different records stored to a VNA may already include original Patient ID values and an MPI value, but one generated by a different MPI system, it is often desirable to have a VNA that can manage multiple patient ID and domain paired values (not limited to a single MPI value).

Limited Storage Technology Options

Many PACS vendors offer a small number of storage technology options with their system, and the storage system often has to be purchased from the PACS vendor—it is often in their interest to limit the configuration variation available to limit their service and support costs. Many do not provide any archiving storage interface capabilities to make use of desirable storage system specific APIs; simply providing a file path method (e.g. NFS or CIFS) instead.

Limited Ability to Manage a Record's Lifecycle

Most PACS assume that all stored records must be kept forever and provide limited to no information lifecycle management capabilities; meaning that the archive will continue to grow as long as new studies are acquired and stored. If a company has more than one PACS from different vendors, the capability to manage the patient's imaging record lifecycle will vary based on which system the record is stored in. Legal allowances and corporate policies may make purging records possible and even desirable.

Limited Metadata Available

As the core of many PACS were designed over a decade ago, some common design choices were made to optimize performance based on the computing power available then. One design pattern found in several PACS¹⁰, is that when archived study images are cleared from cache, some of the more detailed metadata, such as some or all of the series and object metadata, is purged from the database. The rationale is that it improves the database performance by keeping the data stored within it to the minimum required. If data is de-archived, the purged metadata is parsed from the object headers and restored to the database.

An implication of this design is that, while a study's objects exist only in the archive, some of the metadata is not available through the query interface. This can have implications on image display systems, such as Enterprise Viewers as they

⁹ Medical Record Number

¹⁰ Especially those developed some time ago.

often use the detailed series and object information to optimize what specific data to retrieve as a priority. VNA often keep more metadata available in the system, as they are explicitly designed to perform as a data management and sharing platform among multiple systems with varying needs.

Deployment Flexibility

As most PACS servers and storage were physically deployed in the hospital, in a data center or even in the department itself, there are sometimes limitations in how it can be deployed. There are often system design constraints related to the network among the system components. As it may be desirable that some part or all of the VNA may be deployed in the cloud, the system's ability to operate with some components (servers, cache, archive storage, etc.) across a LAN and WAN (with firewall, proxy, etc.) may be important. Often PACS that were designed before this physical deployment model was popular have issues when deployed in this manner.

Integrating a VNA with PACS—What to Consider

So, we have established that a VNA and a PACS server are, in fact, different animals. But they do need to be integrated. And, while DICOM integration can be straightforward, there are a few things one should consider and plan for when connecting a PACS with a VNA.

Infrastructure Proficiency

Network, compute and storage performance will dictate reliability within an integrated VNA and PACS environment. Scalable, flexible and automated solutions should be considered, where full link utilization is available and low latency and maximum availability drives improved end-user experience.

External Archive Capability in PACS

A common design pattern in PACS is to maintain a cache and an archive for storing studies. Cache is used to store recently stored or retrieved studies, and it is typically described by how many “days” of storage it has, but it is generally set up with a fixed size of fast storage (which can be expanded). The archive, meanwhile, continues to grow as more studies are stored—until they are purged by manual action or automated process¹¹. Often DICOM objects for a given study are stored in the archive using some kind of packaged format (like [ZIP](#), [TAR](#), or proprietary [BLOB](#)).

How a PACS deals with using an external system as the archive will vary. Some will support an external DICOM system, like the VNA, as the archive—marking study data in the cache as eligible for deletion once the study is confirmed to be archived to the external system (typically by a DICOM function called C-STORE). Other systems will be capable of performing a C-STORE, but without the capability to treat the study as archived; meaning, it can store the study to the VNA, but the study is not removed from the cache, or it also has to be stored in the PACS internal archive (doubling the storage consumed for the archive).

Prefetching and Relevancy

¹¹ See info on ILM in this paper.

When studies are archived by the PACS to a VNA, and the studies are subsequently deleted from the PACS, they may be retrieved back to the PACS, as needed. There are a couple of common ways that this de-archive of the study is triggered.

Automated Prefetching

The first is when a new exam for the patient is scheduled, an automated process often called prefetching can discover the existence of prior imaging exams (using a C-FIND, for example), and retrieve them (using a C-MOVE, for example). Which exams (see note on relevancy below) and how many exams are retrieved is often configurable.

When retrieving prior imaging exams, it is desirable that ones that are relevant to the newly acquired exam are selected. There are different schools of thought about how to determine what type of exams are relevant, but typically prior exams with the same or similar body part examined, as well as the same or a similar modality type, are considered relevant.

When integrating a PACS to a VNA, the system that will be used to perform the prefetching and relevancy rules needs to be decided. Some things to consider:

Does the PACS include a prefetching and relevancy rule capability? Some systems that do not support an external VNA, and keep all archived studies within its own internal archive, may not have the ability to automatically pull prior exams from the VNA.

If there are multiple different PACS integrated with the VNA, do they all provide the desired level of configurability within the relevancy rules to meet the PACS users' needs?

If the VNA provides the capability to determine relevant prior exams, and the ability to route them to the appropriate PACS (assuming there is more than one connected to the VNA), it may be desirable to configure this logic in the VNA. Taking this approach provides a centralized configuration for the entire system.

If the VNA is receiving study data from multiple systems at multiple facilities, there will typically be additional complexity to be managed. Assuming that the involved facilities use different procedure coding schemes, the VNA may be faced with studies whereby the body part examined in the study is defined differently, making relevancy rules with predictable and reliable results difficult. The use of an ontology or lexicon to normalize the procedure information within the VNA is one approach, but often introduces issues when the data is retrieved by a PACS that is not using the same procedure code sets.

Manual Retrievals

The second method of retrieving prior exams is to manually pull them on demand. While this is quite simple to achieve, the time from request to when the study is displayed in the PACS client is an important aspect to consider—obviously, users will not be happy waiting, even for large studies.

Interoperability of Entire Imaging Record

Interoperability of the full imaging record—including so called evidence documents, such as presentation state information, key image identification, and reports—has been feasible for a number of years.

Here is a quick reference to some common evidence document types and purposes...

IHE Integration Profile	DICOM Object SOP Class	Purpose
Consistent Presentation of Images (CPI)	Greyscale Softcopy Presentation State (GSPS) ¹²	Markups, measurements, display parameters (window width/level, zoom level, flip/rotate, collimation, overlays, etc.)
Key Image Notes (KIN)	Key Object Selection (KOS)	Selection of individual DICOM objects, such as images key for report interpretation, surgical planning, teaching, etc.
Simple Image and Numeric Report (SINR)	Structured Report (SR)	Structured text data including the diagnostic report, modality measurements or CAD markers

All VNA will accept the evidence document DICOM SOP Classes, but not all PACS create these objects. Some will store this type of information using proprietary methods, typically by persisting them in the DICOM Object header; but, in some cases, the information is not included in the objects at all (it is only in the PACS database). Studying the DICOM Conformance Statements and IHE Integration Statements of the VNA, as well as those for every PACS and any other connected system (e.g. Enterprise Viewer), to ensure interoperability of this information, according to standards, is an important planning step prior to integration.

Additional Notes

In addition to studying the DICOM Conformance Statements and IHE Integration Statements of the involved systems, it is recommended that historic data is also evaluated. While the current system may be capable of producing standards-based, interoperable evidence documents, some of the older data stored in the PACS may only have proprietary information, as a prior version of the software was not as capable. Also, some of the older data may have been migrated from a PACS that was previously replaced by the current PACS.

Some VNA may attempt to resolve the lack of standards-based, interoperable markups by using metadata coercion (aka “tag morphing”) to extract and convert the presentation data to the standard form (DICOM GSPS). Caution should be exercised before using this method, however. Over time, a PACS vendor may have changed how they have represented the presentation state information, but this may not be obvious to the VNA vendor. Considering that presentation state information includes measurements (calipers, etc.) and other position critical markups, errors made in the conversion could have patient safety implications. Any created presentation state information not provided by the PACS vendor—such as by the VNA or 3rd party tool—should be carefully validated by qualified imaging experts.

Content Monitoring

Because source systems (RIS, PACS) need to transfer information to the VNA, in such a distributed system, problems will inevitably occur that will cause the source systems and the VNA to not be completely in-sync. For example, a failure to negotiate a particular SOP Class or failure to receive a particular object due to incompatibilities in certain key attributes.

¹² A variation of this SOP Class for color images, Color Softcopy Presentation State (CSPS), is also defined in the DICOM standard, but is far less common than the greyscale version.

Or it can be simply due to a network or queuing delay. Or it can be due to some manual action performed at the source that changes data when the changes are not synchronized to the VNA (e.g. manual quality control operations). As a result, the VNA becomes out of sync from the source(s). Therefore, it is highly desirable to have a content monitoring system that can alert the administrator if the sources and VNA start to drift apart.

Data Synchronization

A common misconception is that DICOM is fixed content; that is, once generated it does not change. In reality, changes to the study metadata (patient info, order/procedure info) are fairly common, as are the addition of new objects (e.g. a new series of images generated from an advanced visualization system, added to the same study). These changes are well-understood and relatively easy to manage. It is when objects within a study are deleted or the structure of the study is altered (some images moving to a new study), that issues arise. Once a study is sent from PACS to the VNA, if an object is deleted from, or the structure altered in, the copy of the study in the PACS, there is no commonly available method for communicating this type of change to the VNA. A common manual procedure is to delete the study from the VNA and to transmit the altered study from the PACS—in a high-volume environment, this can be more than a full time job, depending on the frequency of the changes. IHE [IOCM](#) (Imaging Object Change Management) promises to provide some automation to this operation, but few systems support this integration profile at the time of publication.

The Future

The 'A' (archive) will continue to come out of PACS and the EMR will manage the patient's imaging record, not the PACS. A VNA will increasingly be considered an EMR component, as will the Enterprise Viewer. Both will be purchased and managed by corporate IT, or whomever "owns" the EMR. The storage, virtual servers upon which the VNA runs, and network infrastructure will be the responsibility of corporate IT. Integration with other organization's systems, through an HIE or directly, will be the responsibility of corporate IT. Radiology and similar departments will continue to manage the modalities, the acquisition protocols, and the remaining parts of the PACS, for now.

Analytics and Business Intelligence (BI) that span the complete EMR (and HIE, if it exists), including imaging information in the VNA, will be key to operational and economic improvement. Full text searching will fill gaps where structured and coded information is not available.

The storage of the files in the long term archive will move to the cloud, as it proves reliable. Eventually, the full EMR, and all its subcomponents, will also move to the cloud.

Enterprise ordering will be procured by the EMR team and supplied as managed service to Radiology, just as the RIS is being provided in the EMR in many cases today. Standardized and codified order sets will be mandatory.

In the mid-term future, workflow management will move out of the PACS and into a shared, enterprise business process orchestration platform. This platform will be capable of spanning multiple facilities, enabling resource (modality, physician, equipment, radiologist, etc.) sharing across a broad constituency, delivering significant efficiency improvements. Report creation tools, such as the voice recognition systems employed in Radiology today, will be in the EMR.

The enterprise viewing system, integrated with the EMR and managed by corporate IT, will continue to evolve, replacing dedicated PACS workstations and advanced visualization systems as they do.

Experienced Radiology IT and informatics professionals will be sought after as technical imaging subject matter experts in managing, configuring, and supporting these enterprise systems for a broad range of consumers.

Additional Reading

PACS in 2018: An Autopsy

<http://link.springer.com/article/10.1007/s10278-013-9660-1/fulltext.html>

Informatics Challenges—Lossy Compression in Medical Imaging

<http://link.springer.com/article/10.1007/s10278-014-9693-0/fulltext.html>



About the Author

Don Dennison has worked in the imaging informatics industry for over 13 years and is currently serving as a consultant. Previously during his employment with a PACS vendor, he oversaw a team focused on interoperable, longitudinal patient imaging records within electronic medical record systems and across enterprises.

Don has always been interested in innovative technologies. He has managed product development teams from 1 to 150, defining a vision and creating new products that change the way IT is used to solve real challenges within healthcare.

Some highlights of Don's professional experience include:

- Chair (SIIM 2014) - "Web Technologies: The New Healthcare IT Standard (Hot Topic)"
- Author (Journal of Digital Imaging) - "PACS 2018: An Autopsy" (Dennison, JDI, 2013, 10.1007/s10278-013-9660-1)
- Author (Journal of Digital Imaging) - "Informatics Challenges—Lossy Compression in Medical Imaging" (Dennison, Ho, JDI, 2014, 10.1007/s10278-014-9693)
- Author (Journal of Digital Imaging) - "Where to Build It" (Dennison, JDI, 2013, 10.1007/s10278-013-9658-8)
- Panelist (SIIM 2011) - "Applied Learning Vendor Tie-in Session: Image Sharing: A Vendor Perspective"
- Panelist (Canada 3.0, May-2011; MaRS Global Leadership Series, Jun-2011): "Social Networks and Health Care"
- Presenter (Jan-2011) - "Images, Electronic Health Records, and Meaningful Use: A Vision for the Future"; sponsored by National Institute of Biomedical Imaging and Bioengineering (NIBIB) in a talk entitled "Integrating Words & Pixels: The Value of Images in the EHR"
- Co-author and Presenter (SIIM 2010) - "Evaluation of the Two Major Strategies Often Proposed for Sharing Imaging Information across Multiple Facilities (e.g., HIE, RHIO) and Delivering Images to the EMR Using IHE"
- Presenter (HIMSS 2009) - Interoperability Showcase - "Using IHE to Enable Medical Imaging Archiving & Visualization Across Your Enterprise"
- Panelist (SIIM 2008) - Closing General Session - "Building Bridges: The Challenges Facing Vendors and Imaging Informatics Experts"

All Original Content © 2013-2014 Don K Dennison Solutions Inc., All rights Reserved



Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com