What is a VNA, anyway?

By Herman Oosterwijk, President, OTech Inc.

The term *Vendor Neutral Archive* (VNA) is used by many vendors of PACS software to identify the core of their product--the component that is used to reliably manage, store, retrieve, and query medical images and related information. Unfortunately, there is no "official" definition of the functionality of a VNA. As such, VNA offerings vary from vendor to vendor. The result is confusion among users due to the lack of a baseline that permits a direct comparison of VNA functionality and features.

This white paper proposes a definition of a VNA that identifies different architecture levels and their specific functionality. The various components which make up the core of a healthcare image and information management system will be discussed. It also provides the key advantages of a VNA compared with a proprietary solution. A checklist which can be used when evaluating and/or requesting a RFI/RFP for a VNA is included as an attachment.

It is somewhat unfortunate that the term VNA has stuck as it describes not really what a VNA is supposed to achieve. Ideally, a VNA is an archive developed on an open architecture that can be easily migrated and/or ported to interface to another vendor's viewing, acquisition, and workflow engine to manage medical images and related information. Instead of "Vendor Neutral", a better term would have been "Architecture Neutral", "PACS Neutral," "Content Neutral" or "Third-Party Neutral." Keep in mind that each commercial archive is provided by a specific vendor (we are not considering open-source offerings, which may be the only truly vendor neutral solution). Therefore, the word *neutral* does not really make sense. However, the VNA acronym has become fixed in both the medical IT user and vendor culture, so we will not make any attempt to re-phrase the term.

What is the problem that a VNA intends to solve?

As of today, there are literally thousands of PACS installed and in operation worldwide. Most of these are working satisfactorily, and provide a more efficient and effective solution to manage medical images than was achieved with analog (film-based) solutions. However, PACS users are experiencing growth pains due to two issues: The first one is caused by needing to serve not just the radiology department but also other departments and specialties. These can include departments such as cardiology, radiation therapy, dentistry, as well as other "ologies" such as ophthalmology, endoscopy, and other specialties that use visible-light acquisition modalities. This expansion can be expected to see exponential growth when pathology converts its exams to digital images. As more departments seek access to PACS, the proper identification of patients and exams becomes an issue. Often, each department will have its own patient and exam requisition system, Accession Numbers may or may not be used for identification, and patient and exam IDs may be assigned by each department. There are also very different performance and storage requirements for each of the modalities in each of the departments.

The second PACS growth pain being experienced by users is when they change PACS vendors. This has proven to be a costly, lengthy effort for early adopters. In addition, important information such as image annotations, Key Images, and changes to studies are sometimes lost.

These are the major issues that a VNA addresses: Scalability beyond a department-level archive solution that delivers true data integrity for image data and related information; the capability to change PACS vendors without the need to migrate or convert image data and to create true independence from a specific storage solution, i.e., the hardware platform that a PACS vendor might force upon a prospective customer. This changes the role of the traditional PACS archive to one of a transient, easy replaceable component that stores only those images that need to be available for a limited amount of time (e.g., 3-6 months), while the VNA provides long term, vendor neutral image and information management.

Before we go into the details of what a VNA is, it might be easier to define what a VNA is NOT, as many offerings on the market claim to offer a true VNA while not really meeting all VNA requirements. In order to do this, we propose five diverse archive architectures, identified by 5 different levels. Each level has increasing functionality, with only Level 5 achieving VNA status:

- Level 1: This is the traditional DICOM archive, which represents most of the installed base as of today. It serves typically only one department such as radiology and/or specialty such as ultrasound, digital mammography, nuclear medicine, or radiation therapy. All information is stored in the DICOM format, documents are converted to DICOM Secondary Capture and/or encapsulated PDF's. Even Waveforms are in DICOM formats, such as is used for cardiology. Studies are characterized by being scheduled by a single scheduling system issuing a unique Patient Identifier and orders are identified by a unique Accession Number.
- Level 2: Multi-departmental archive, which typically can serve multiple imaging departments and/or specialties, such as cardiology and dentistry. It can manage also some limited non-DICOM information such as physiological data and EKG's for cardiology, JPEG's for dentistry, ophthalmology, or endoscopy. Patients could have multiple Patient ID's and/or scheduled by multiple schedulers resulting in multiple, non-unique Accession Numbers.
- Level 3: Multimedia archive, which expands the capability of objects to be archived, such as MPEG's for endoscopy and speech pathology, as well as a wide variety of text documents (PDF, Word, etc.). This is the level where digital pathology slides are managed.
- Level 4: Enterprise archive, which serves as the backbone for multiple institutions, hospitals, and/or regions. A fully functional Master Patient Index (MPI) is needed to reconcile patients.
- Level 5: A VNA, which provides a patient-centered, open architecture that allows for managing the image. Related information can be exchanged using XDS protocols and patient demographics be reconciled using PIX/PDQ protocols.

As shown in figure 1, each level has an increasing scope/domain as well as format/standards that are supported.

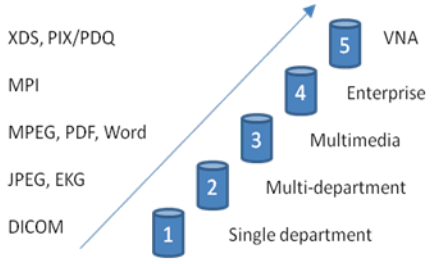

Figure 1: relationship between domain, standards and corresponding architecture level

One of the key elements that is not shown in the figure is the implicit paradigm shift that is taking place when moving from Level 1 to 5 with regard to data ownership. Most DICOM archives are managed by the vendor, which is contrary to common IT practices where the internal IT department has full access to the data and database. The majority of the DICOM archives are locked up by the vendor. PACS administrators, in many cases, are not even able to get system access privileges. This means that simple queries, changes, and updates always have to be performed by the vendor. Unless negotiated upfront, users may not even have the archive's database schema specification. It is generally at Level 4 that data ownership is typically moved to the user.

Given the historic perspective above, the proposed definition of a Level 5 archive, or VNA, is as follows:

- ***A Vendor Neutral Archive (VNA) is a medical device that provides scalable image and information and life cycle management so that images and related information can be queried, stored, and retrieved in a manner that is defined by open standards at multiple department, enterprise, and regional level while maintaining patient privacy and security. Characteristic for a VNA is that it provides a patient-centric approach that transcends upgrades and changes of the different viewing, acquisition, and workflow management components as they should be interchangeable without having to migrate, convert, or change the data formats or interface of the VNA.***

Let's discuss the requirements of a VNA using the underlined keywords in the definition:

- A VNA is a Medical Device: A VNA is a device that meets the U.S. federal requirements of a Class 1 Medical Device. This means that the VNA manufacturer registers its product and device manufacturing location with the FDA, makes sure that the device (including its software) is developed according to a quality system (aka Good Manufacturing Practices as defined by U.S. Regulation 21 CFR 820), and has a complaint handling procedure in place. This requirement differentiates "medical grade" devices from off-the-shelf consumer devices.

- A VNA is Scalable: It is hard to exactly estimate a facility's required storage capacity. Most users seem to underestimate what it required, although some of them over-estimate and do not make good use of their investment. Factors that impact a growth in storage utilization are:

    - Purchasing new modalities: The increasing popularity of the latest generation of multi-slice CT scanners has had a big impact on the amount of data generated.
    - New imaging protocols: A new MR protocol might suddenly generate thousands instead of hundreds of images.
    - Equipment upgrades: Upgrading a digital room with a high-definition imaging subsystem, from a 512 to 1,024 acquisition matrix, would quadruple the storage requirement.
    - New departments are getting on-line: Most administrators are wary when cardiology comes on-line because of that specialty's generation of large cine runs consisting of thousands of images. Digital mammography also has a major impact on storage capacity; a typical mammo exam of four images could be 160 MB in size.
    - New specialties: Everyone is apprehensive about the storage requirements of digital pathology. Their data generation is much larger than of any current imaging modality. i.e. in the multi-gigabyte range/week.

    In general, scalability is a difficult and somewhat vague requirement, as the upper limit of the scale and also the corresponding performance requirements are hard to define. For example, imagine that the VNA supports the storage of all images taken by 100-facility chain of hospitals . Another good example would be all images of patients residing in a certain metropolitan area or a state. One can imagine that this database could easily encompass many millions of people, with the amount of images and related information a multiple of that; therefore, it is hard to specify an upper limit.

    With regard to performance, in many cases, a new procedure is scheduled in advance and required comparison studies can be pre-fetched to local cache and be made available by the local PACS within a few seconds after selecting that study. However, in the case of an unscheduled procedure, where the information has to be retrieved real time from the VNA archive, the retrieval time could be 5-10 seconds, or, up to one minute for large images such as a Cardiology run. The good news is that these occurrences should be infrequent and therefore will not have a major impact on productivity.

    Another key to scalability is that the information can be preserved and should not have to be migrated and/or ported in case the data exceeds certain semi-artificial limits, such as 1 million images. Note that this requirement does not mean that the database might not have to be ported as this might be needed to facilitate a larger number of database records.

- A VNA provides Image and Information Management: The Image and information management component of the VNA is provided by a database. The image manager or database stores a subset of the Attributes available in the DICOM header of the objects with additional metadata needed to manage the images, such as a pointer of the location of the images, date of last retrieval, and so on.

Two critical features are needed, the first being a mechanism to perform statistics and troubleshooting. For example, a user should be able to track the number of images stored by modality by date so that a trend analysis can be performed. A user who is not monitoring their archive utilization on a regular basis is setting themselves up for potential surprises.

In addition, QA/QC tools are needed to track and log transactions to find "lost" images that were improperly identified and end up associated to the incorrect patient, study, or other record, and subsequently fix these. Vendors also need to provide system administrator tools for the archive to support management and allow for database maintenance, such as records to be cleaned and modified. This is needed because information that identifies the images is at some point entered by a user, which is prone to error. In other cases, information will need to be modified after an imaging exam, particularly in trauma cases when the technologist might be focusing on getting the exam done prior to making sure it is properly scheduled and all the information is correct. In this case, a system administrator might get a request to retrieve a study performed at a certain time at a particular modality belonging to "John Doe" and correct the information.

Changes and updates in patient demographics have to be exchanged between the PACS archive, VNA, and external image management devices such as the department or institution information system.

- A VNA includes life-cycle management: A VNA needs to deal with image deletion and retention aka as lifecycle management. Some images might have to be deleted, e.g. within 30 days of the study date that were of non-diagnostic value. It is not sufficient to "flag" a database entry to identify deletion (as is common with some of the major PACS vendors), as this information is again proprietary, but, rather that the images be physically deleted.

  Images and related information might have to be deleted as the legal retention period for keeping information expires, typically 7 years for adults. The retention period depends on body part (e.g, for mammography there could be a longer retention period), patient age (e.g., keep pediatric images until the patient reaches 21 years of age), and study type. It also could depend on institution policies as some institutions might decide to keep studies longer for research purposes. The retention period also depends on whether, or when, a study was accessed recently. If a study does not need to be retained after a certain period, but based on retrieval records, it is shown to be accessed on a regular basis, deletion might not be warranted.

- A VNA manages images as well as related information: There are basically two types of information objects to be managed; the majority, at least for now, is based on the DICOM format, and the second type is the non-DICOM objects. The DICOM objects are images, as well as Structured Reports (SR) and Presentation States (PR). The Structured Reports are typically CAD data or measurements such as performed during ultrasound exams, but could also determine which images are identified as so-called "key-images" by the interpreting physician. These key images allow for subsequent retrieval to only pull the most important images of a study, or "key" objects. The DICOM Presentation State objects are also critical as they allow for a vendor-neutral representation of overlays, shutters, and other user manipulations (such as zoom) so that the presentation can be preserved. The SR and PR objects are identified with the same study identifier as the images they relate to, enabling them to appear in a study-based query. Other DICOM objects such as defined for radiation therapy could be managed as well. Non-DICOM objects include waveform documents, PDF, XML, other text-based formats, TIFF, multimedia such as MPEG formatted files and even raw data as generated by specific modalities (e.g. spectroscopy data from an MRI).

- A VNA provides the capability for query, storage, and retrieval: A VNA implements the so-called image manager and image archive actors as defined by Integrating the Healthcare Enterprise (IHE). The image manager is a database which provides responses to query requests. When a data object is sent to the VNA for the first time, important information from the header is captured,

potentially cleaned, and stored for queries with additional meta-data needed for the image manager. Cleaning could consist of capitalization, to prevent issues with case sensitivity (e.g., “SMITH” does not match “smith”), unless this case-insensitivity is an inherent feature of the database that is used. There are also typical rules to limit the number of responses to prevent users from doing “wide-open” queries which might impact the database performance (e.g., limit the number of maximum responses by default to 100).

The archive provides secure storage and retrieval as well as routing and prefetching capability. Prefetching is typically performed when order information is received resulting in sending prior exams to a specific PACS and/or workstation. The selection of the prior exams depends on configurable parameters such as body part, modality, study date, and study/series description. Routing is mostly used to allow for back-up and/or mirrored copies of the information or for Teleradiology purposes.

- A VNA supports open standards: The communication with the VNA takes place using the DICOM Storage Service protocol, while the searches are supported by the DICOM Query/Retrieve Service protocol. There is a list of required (R) and unique (U) attributes, or keys, defined in the DICOM standard that has to be supported by every DICOM-compliant archive (see table).

Patient’s Name	(0010,0010)	R
Patient ID	(0010,0020)	U
Study Date	(0008,0020)	R
Study Time	(0008,0030)	R
Accession Number	(0008,0050)	R
Study ID	(0020,0010)	R
Study Instance UID	(0020,000D)	U
Modality	(0008,0060)	R
Series Number	(0020,0011)	R
Series Instance UID	(0020,000E)	U
Instance Number	(0020,0013)	R
SOP Instance UID	(0008,0018)	U

The DICOM standard is almost exclusively used for the communication between modalities and the PACS, and the PACS and VNA. In addition, many vendors have implemented database queries from their workstations, typically using SQL, into their database in order to provide greater efficiency and functionality. A SQL interface is also important for system administrative tools for statistic reporting and troubleshooting. Standard exchanges with information systems also are DICOM- and/or HL7-compliant to exchange patient and study updates.

The image manager and archive interface supports DICOM Modality Performed Procedure Step and Storage Commitment as defined by the IHE scheduled workflow profile. IHE also specifies the exchange of documents and images with external information sources using the XDS and XDS-I profiles.

With regard to the internal data format storage, there is unfortunately no archiving standard defined to address this as the standards typically deal with the interface only. A de-facto archiving standard for image and related information has become the DICOM exchange media standard also known as “Part-10”, without the DICOMDIR structure (which is meant to be used when archiving images on CD or other external media). The so-called metafile (“Group 0002 Attributes”) are important as it defines the encoding of the DICOM objects (e.g., JPEG, MPEG, wavelet, etc.). The Part-10 requirement preserves the DICOM object, which in most cases are pixels, and the DICOM header information. Some formats, such as PDF and MPEG, can be encapsulated as DICOM objects, depending on how they are received, but also should be able to be archived in their native non-DICOM format in order to meet the requirements of a true multimedia archive. Additional metadata might be stored in such a way as to include the image or object location, last retrieval, and whatever an archive needs to manage the object.

- A VNA supports multiple departments, enterprise and regional architectures: Centralizing storage management is definitely much more cost-effective than doing it on a department level. Especially the back-up, high availability features and business continuity is easier and more efficient to provide when the storage solutions are centralized. The same advice applies for the physical storage: having a computer room in radiology and cardiology and pathology does not make sense if there is already a very capable IT staff that has secure and reliable archive space available at the enterprise level. One of the practical arguments from department managers and/or system administrators against having an archive supported by the IT department is that one has to relinquish control and does not have always full access and privileges to manage and maintain the system. A good compromise would be a separate room in an institution's IT center that is accessible by both the IT and radiology support staff. In any case, enterprise-wide storage is definitely a trend that will continue to evolve, and, in many cases, regional archiving makes sense as well, especially for hospitals that are affiliated or owned by the same organization. Critical for the support of multi-department support is the reconciliation of multiple patient identifiers via support for a Master Patient Index (MPI). Many traditional DICOM archives rely on unique Accession Numbers as well, which is something that cannot be guaranteed at the multi-department, enterprise, or regional level. Patient demographics can be reconciled using the PIX/PDQ capabilities as defined by the IHE.

- A VNA maintains patient privacy and security: A VNA needs to support audit trails, such as a log listing of who is accessing what information and when, in a manner that is easy accessible and retrievable. Authorization rules are typically provided by a different application but could also be part of the VNA. This also satisfies legal concerns, as it can be proven that a certain physician retrieved a particular case. Instead of each device keeping these audit trails in a semi-proprietary manner, it makes the life of a system administrator much easier if this information is updated and maintained in a central location, in a standard format as specified by IHE in its ATNA profile definition. Data integrity is also critical, which is why some vendors implement digital signatures to make sure that alterations can be detected. Lastly, there are some institutions that require non-volatile media for their long-term archiving such as magneto-optical devices (MOD), or other write-once media. The reason is that the information on rewritable media such as RAID's could be compromised by viruses or other malicious software, which is not possible if the data is burned-in to a medium that is not erasable or cannot be modified.

- A VNA is patient centric: The department-centric model has served early PACS implementations; however, it fails terribly when moving to a patient-centric model. This is when information from different departments has to be provided in a multimedia format (i.e. some of it DICOM, some of it not) for a single patient. For example, most vendors who provide so-called "integrated" radiology/cardiology solutions still have a separate database and image storage for each specialty and implement a middleware application that merges retrievals from both databases. One can imagine the impact when yet another department goes digital and needs to be integrated. A similar issue arises when access needs to be provided from more than one PACS. These tight connections between the multiple components is the most visible for a teleradiology company that serves different sites from multiple vendors with a different workstation for each vendor in order to access images. A patient-centric solution rather than department and/or vendor approach provides a single user interface and single-access capability.

- A VNA transcends upgrades and changes of PACS and allows for PACS to be interchangeable: Most PACS vendors have a very tight, non-standard connection between the PACS core components, especially when it relates to the interfaces to the archive and image manager/database and workstations. In addition, the data formats could be archived in a proprietary manner, on the assumption that all data coming into the archive is DICOM based. This worked for the initial years of PACS implementations, which were department centric; however, this does not meet the requirements for true enterprise storage and does not allow for PACS vendors to be changed without a lot of effort and expense. A true VNA allows a PACS vendor to be changed with relatively few interface changes by the new PACS vendor.

- A VNA does not require data migration and/or conversion of data formats: Data migration, when exchanging PACS vendors, has become a major expense, requiring significant resources and delaying new PACS implementations. Migrating data to a format that might have to be migrated again in a couple of years does not make sense; a much better choice would be to migrate it only once to an open and widely supported standard.

  This does not mean that the actual objects would stay on the same physical media for ever. As hardware storage technology improves, data might be moved (let's say every 5 years); however, the data formats would not change and there would be no disruption of the image availability. The same applies for storage software technologies, improvements to support different database technologies, and virtual storage architectures. Again, the latter should not impact the image availability and might be done to improve performance and/or cost rather than be a forced, costly migration because of vendor changes.

- A VNA does not require interfaces to be changed: The VNA has plug-ins that allow different applications and/or vendors to be connected. There should not be a need to custom-develop interfaces to various vendors as PACS vendors should support the various standards. However, some vendors have peculiar implementations and might require specific information in the DICOM headers that makes their workflow software function optimally. An example would be the definition of specific study and/or series descriptions to match the prefetching of corresponding studies and selection of specified hanging protocols. In the absence of a true vendor neutral interface at the PACS, middleware might need to be used to adopt an interface to mitigate a vendor's idiosyncrasies.

There are other requirements that might be part of the VNA; however, by themselves they do not create a true VNA, such as:

- SSP and ASP capability: The archiving and archive-management component is typically the most expensive element in a digital healthcare imaging system purchase. Most vendors offer a Storage Service Provider (SSP) or Application Storage Provider (ASP) model whereby a storage and retrieval fee is charged by the exam or by the retrieval. When using an ASP model for data storage, one need to make sure there are proper guarantees in place in case something happens with the ASP provider, i.e., that the data originator maintains ownership of its data. When taking the proper precautions, this model can eliminate a substantial upfront fee, and it is expected that for an increasing percentage of the users, this will be a sensible model--especially when in-house IT skills and expertise are somewhat weak as facilities outsource data management. An SSP/ASP provider could offer a VNA, as a matter of fact, most service providers do offer VNA capability; however, a SSP/ASP is not always a VNA unless it meets all of the requirements listed above.

- Cloud or Grid storage: Cloud or Grid storage are examples of virtual storage solutions, whereby the actual storage is virtualized to the point that a sender or retriever has an access point that allows for data exchange without knowing where the data resides. The cloud or grid takes care of meeting certain performance requirements, back up, and high availability through data duplication and other techniques.

- Uni-viewer access: In addition to the radiology workstation, most PACS vendors provide a plug-in viewer that allows for images to be displayed, for example, as part of an EMR portal for a physician. This is a temporary solution and only provides access to the images in the PACS archive and database. A better solution is that the VNA provides access to its multimedia information source using a universal viewer. This should be provided using a standard interface such as WADO for DICOM objects.

- Tag Morphing: Changing the DICOM header information is sometimes referred to as "tag-morphing". There are three reasons for making changes in the DICOM header:

    1. Neutralization: Changing any of the incoming DICOM header information to clean up the data and "standardize" it, especially when received from modalities that are not quite DICOM compliant. An example could be the presence of an invalid UID, a study description that is not standardized, institution name to be added, and so on.
    2. Identification: Transfer the incoming data to the proper patient identification domain by adding MPI identifiers, Accession Numbers, and/or changing the ID's.
    3. Customization: Changes in the DICOM header for outgoing data might be needed to adopt images and related information to quirks and peculiarities of certain vendors, such as to make workflow engines and hanging protocols work.

    Some authors and/or vendors refer to this process as "bi-directional dynamic" tag morphing, referring to making the changes in the input as well as output of the DICOM header and the fact that this happens upon demand (dynamically). One could argue that this tag-morphing should be done by the PACS vendors (perhaps through middleware) rather than the VNA vendor; however, in some cases, the VNA vendor might implement this just to ensure the system is working.

- SOP Class/Xfer Syntax conversion: This feature is similar to tag morphing as it converts the information based on the capabilities of the receiver; however, it is more drastic. It changes the SOP Class, such as converting a new enhanced MR or CT multiframe object into a traditional MR or CT object, or an encapsulated PDF into a Secondary Capture. It also might be able to convert the transfer syntax, such as JPEG or Wavelet encoding, into an uncompressed syntax, or the other way around to convert an uncompressed image into a JPEG compressed image. This feature could be quite sophisticated depending on the range of SOP Class conversions and different compression types supported.

- Disaster Recovery, high availability and back-up: The requirement for disaster recovery and back-up has always been good practice; however, with the U.S. federal HIPAA regulations they are no longer an option but a requirement. Disaster recovery and high availability are two separate requirements, although they both achieve the same result: Providing storage and retrieval of the information reliably, independent of hardware and software failures, or disasters. Disasters are typically from the outside, and might include lightning, floods, extreme weather, or even terrorist actions. Something which is often overlooked is simply the physical location of the devices, for example, some devices are located in the basement of a facility, in an area that is prone to flooding.

    High availability is typically used to identify the recovery from hardware and software failures. There is high availability on the local physical level through the use of RAIDs. These have the capability to recover from a lost drive by using parity information stored in several other disks. Most of these disks are "hot-swappable," which means they can be replaced without having to bring the system down. However, using RAID technology only is not sufficient. Recently, a hospital in Oregon lost more than 5,000 x-ray studies due to the simultaneous crash of multiple RAID's. Because of the loss, the hospital had to issue letters to 900 patients informing them of the situation along with an offer to redo their exams at no charge.

    However, there could also be a network, infrastructure, or software issue that brings down the complete unit, in which case another unit should be available immediately or within short notice to take over the archive function. To provide this high availability, there are several approaches: One can write the information to two identical archives (mirrored approach), which means that there is always a duplicate copy accessible. It is even better if one were to locate the second archive at least 30 miles from its original, which would not only provide protection against hardware/software failure but also against local disasters. Mirroring the data is not sufficient, there also needs to be a back-up strategy. The reason is that the database could become

corrupted through malicious software (virus, malware, etc.), or a manufacturer software upgrade or maintenance causes the corruption of the database.

Last but not least, not every PACS vendor is ready to connect to a VNA as there are certain requirements that their system has to meet, in particular:

- Storage Commitment SCU support: The PACS archive needs to be able to "hand-off" the responsibility for storage to the VNA which requires it to initiate a Storage Commitment request, i.e. support DICOM Storage Commitment as a SCU. Unfortunately most PACS only implement SCP, assuming that their archive would be the final and long-term storage provider.
- Patient and Study Synchronization: If there are any changes in patient demographics, the study information has to be exchanged between the PACS and the VNA.
- Status Updates: A PACS typically has a workflow manager which keeps track of the study status, such as whether it is read/completed or unread. Pre-fetched studies from the VNA have to be identified as "completed" or "read" or they could end up on the unread worklist.
- Life Cycle Management: Image deletions have to be coordinated between the PACS and VNA, and there should be an aging process for images to be deleted from the PACS after a configurable amount of time (e.g., 6 months).

More detailed requirements can be found in the appendix "RFI/RFP checklist".

Conclusion:

The term *VNA* is widely used by healthcare image and information management vendors to describe their archive offering. Unfortunately, the term is very loosely defined and requirements to meet VNA criteria are not widely agreed upon and accepted. In an effort to achieve consensus about the definition and characteristics of a VNA, a 5-Level archive model is proposed with increasing domain coverage and standards support with accompanying requirements.

Acknowledgement:

This white paper was supported by an educational grant from TeraMedica. For more information on TeraMedica's VNA offering, please go to www.teramedica.com.

The author encourages and appreciates feedback and can be reached at herman@otechimg.com.

Appendix: VNA RFI/RFP Checklist

	Yes	No	Future Date
Medical Device Certifications:			
The device is registered with the FDA as a Class 1 Medical Device, registration number: (provide)_____, or alternatively registered as part of a 510(k) filing, provide registration (K) number: _____			
Vendor manufacturing/development facility is ISO certified			
Scalability:			
The device must be scalable (i.e., support queries and retrievals without noticeable performance degradation) when the number of studies increases.			
Database/Image Manager:			
The vendor shall provide user friendly access to usage statistics which includes--but is not limited to--storage and retrieval rate based on institution, department, modality, and study description			
System administrator tools must allow for patient and study merge			
Tools should allow for changing any DICOM Attribute in the image header and database			
Query tools (SQL is acceptable) shall be provided to access the database for any searches that are needed to locate “lost exams” and perform any analysis (based on images stored) that an institution sees fit			
Changes in image headers should be able to propagate to all images in the respective Exam, Series, Study, and Patient level			
Life Cycle Management:			
It shall be possible for images and related information to be deleted by a PACS administrator			
Images shall be able to be deleted based on configurable retention rules which shall include modality type, patient age, study date, and date of last retrieval			
DICOM Image Object Support:			
Storage and retrieval of Enhanced DICOM objects such as--but not limited to--the new multiframe MRI, CT, XA, and RF shall be supported			
A complete list of all supported Image Storage SOP Classes shall be provided			
The device shall be able to support dynamic conversion of new, Enhanced SOP Classes to the traditional SOP Classes based on the destination AE Title			
DICOM Non-Image Object Support:			
Changes in window width/level, zoom, and pan shall be stored and retrieved as DICOM Presentation States			
Overlays, such as those that contain measurements and notes, as well as markers and shutters, shall be stored and retrieved as DICOM Presentation States			
Key images shall be stored and retrieved as DICOM Structured Reports			
Structured reports--such as to store CAD and measurements--shall be able to be stored and retrieved			
A complete list of supported Non-Image DICOM SOP Classes shall be provided			

Non-DICOM Object Support:			
Query, storage, and retrieval of multimedia formats such as--but not limited to--MPEG, JPEG, TIFF, .DOC, .TXT, .PDF, and .XML shall be supported			
A complete list of supported file formats shall be provided			
Database Query:			
All DICOM Unique and Required Query Attributes shall be supported			
A list of all optional DICOM Query keys for both Patient and Study Level Query shall be provided			
A SQL interface shall be provided that shall allow administrator access to the institution's PACS to perform queries			
The number of responses to a “wide-open query” shall be configurable			
The device shall support WADO for retrieving images and related information over a Web interface			
Frame level query should be supported to be able to retrieve one or more frames from a multi-frame DICOM object			
Database Specifications:			
A current copy of the database scheme and tables shall be provided			
All manufacturer specific Attributes in the database shall be defined according to the DICOM specification rules (i.e., Name, VR, VM, etc.) and explanation of its usage shall be provided			
All “clean-up” rules that involve changing Attributes from the header into the database shall be specified, including--but not limited to--capitalization, removal of control characters, and/or punctuation marks			
Routing and Prefetching:			
Autorouting rules shall be configurable to include Institution, AE Title, Station Name, Modality, Date and Time as well as Referring Physician			
Routing shall take place simultaneously to multiple destinations, which are configurable but shall not be less then 4			
Prefetching rules shall be based on HL7 transactions and be configurable to include--as a minimum--Body Part, Modality, Institution, Study Date			
DICOM Image Management Standard Support:			
The device shall support DICOM Storage Commitment as a SCU and SCP			
The device shall support DICOM MPPS as a SCU and SCP			
Supported IHE Profiles: The device shall support the following profiles:			
SWF: Scheduled Workflow for radiology			
PIR: Patient Information Reconciliation for radiology			
XDS-B: Cross-Enterprise Document Sharing			
XDS-I: Cross-Enterprise Image Sharing			
PIX/PDQ: Patient Identifier Exchange and Query			
ATNA: Audit Trails Node Authentication			
EUA: Authentication			
Transfer Syntax Support:			
The device shall support--as a minimum--Big Endian, Little Endian, Explicit and Implicit VR, RLE, JPEG lossy and JPEG lossless compression, and store images in their native (received) format			
The device shall support standard (DICOM) wavelet lossy and lossless			

compression			
The device shall be able to convert (based on configuration) by source and destination (AE-Title), any transfer syntax to any transfer syntax			
When converting Implicit VR to Explicit VR transfer syntax, all private attributes shall be properly identified in the VR as “UN”			
Archive Format:			
DICOM objects shall be stored in the archive in their native format, which includes the metafile information as specified for DICOM media exchange (“Group 0002”), also known as “Part-10” files			
Additional metafile information stored with the information objects shall be specified by the vendor			
Error Handling:			
Improperly identified images, missing DICOM Type-1 and Type-2 Attributes, duplicate Patient ID’s, and/or invalid or duplicate UID’s, shall be flagged as such and stored in a temporary cache that shall be able to be reviewed by a System Administrator			
Patient Privacy and Security			
Audit trails shall include--as a minimum--the access by who and when, shall be maintained and be searchable, and shall be presented in a user-friendly format			
Audit trails shall be stored in a standard format according to the ATNA IHE profile			
A list of allowed IP addresses and AE titles shall be configurable to provide a minimum level of a access security			
If applicable (when allowing access through a uni-viewer), users shall be authenticated			
If applicable (when allowing access through a uni-viewer), single sign-on shall be provided using CCOW			
ASP/SSP Model:			
Vendor may provide (optionally) a turn-key solution (hardware and software), software only solution, and ASP model which is priced on a per-study basis			
Hardware/Software Architecture:			
Vendor shall specify the architecture options (grid, cloud, etc.)			
High Availability/Disaster Recovery and Back-Up:			
Vendor shall explain the high availability, disaster recovery and back-up options such as provided by mirrored and/or duplicated interfaces, hardware and software			
Tag Morphing:			
Vendor shall specify its tag morphing capability--whether it is static, dynamic, bi-directional--and a list of attributes and their values that can be configured			
Vendor shall provide a list of at least three institutions which successfully interfaced with the requesting institution's vendor and software version/release (Vendor ABC, release xyz)			

