

THE BRIDGE FROM PACS TO VNA: SCALE-OUT STORAGE

Authored by Michael Gray of Gray Consulting

Abstract: Moving to a vendor-neutral archive (VNA) for image archival, retrieval, and management requires a phased storage approach due to the capital and operational expenditures involved. The EMC Isilon scale-out approach provides a simple, predictable, and manageable path from PACS (Picture Archiving and Communications System) to VNA.

April 2012

Copyright © 2012 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All trademarks used herein are the property of their respective owners.

Part Number H10699

Table of contents

Introduction	4
Phased evolution toward a VNA.....	6
The bridge to VNA: scale-out storage	8
Your partner in this evolution: EMC Isilon.....	13
Conclusion	14

Introduction

Medical imaging professionals continue to embrace the digital domain. Following the path paved by leaders in radiology and cardiology departments, divisions such as endoscopy, dental, ophthalmology, and pathology are now showing interest in the digital image management systems known as Picture Archiving and Communications Systems (PACS). Medical departments such as oncology and surgery that create video or still-frame images during the course of patient treatment are also looking to implement systems to manage their data. The healthcare industry recently announced that Meaningful Use criteria for Stage 2 will now include an option for medical images¹—a development that will no doubt strengthen interest in imaging and accelerate the deployment of new departmental PACS.

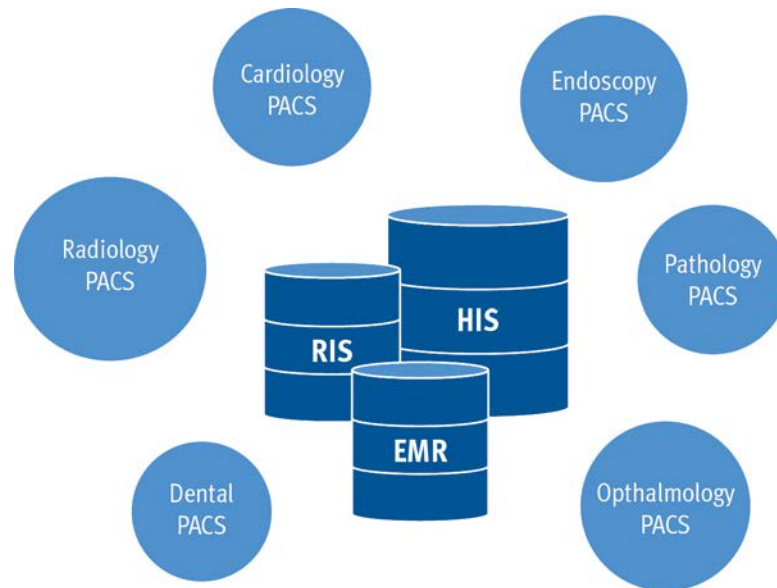


Figure 1. Islands of Picture Archiving and Communications System (PACS) for Healthcare Information Systems

Departmental PACS traditionally come packaged with their own closely managed storage solution, with each vendor offering limited choices. Multiple departmental PACS sharing a consolidated storage solution is a rare occurrence, even if all PACS were provided by the same vendor. One could imagine an organizational landscape three years out with six or more departmental PACS maintaining individual dedicated and different storage solutions. Without thoughtful planning, these new PACS implementations could easily become an IT nightmare.

¹ <http://www.dotmed.com/news/story/18176>

Individual islands of storage are inefficient to manage, requiring different skill sets and experience, and therefore additional staffing. Block-based storage solutions are also largely underutilized, typically requiring 50 percent of the capacity drawing power while storing nothing. Storing all of a healthcare organization's image data in individual department PACS—and, by extension, their dedicated storage solutions—is also an impediment to efficient image sharing. Achieving Meaningful Use of images scattered across multiple disparate platforms presents a significant challenge.

The ideal approach to managing all of the organization's image data is the relatively new concept of the vendor-neutral archive (VNA). The VNA is defined as an Enterprise-class data management system that consolidates primarily medical image data from multiple imaging departments into a master directory and corresponding consolidated storage solution, thus replacing the individual archives associated with departmental PACS. As a consequence, the VNA effectively becomes the unified image data repository for the electronic medical record (EMR) system.

The VNA offers the following major improvements over the separate image data repositories represented by today's departmental PACS:

- Consolidated storage, thus simplifying expansion, upgrades, management, and support
- Ability to accept and manage non-DICOM (Digital Imaging and Communications in Medicine standard) images as well as non-image data
- Normalization of the DICOM headers, thus facilitating data exchange between PACS and elimination of future data migrations required by the replacement of a PACS
- Introduction of sophisticated Information Lifecycle Management (ILM) driven by clinical metadata associated with the images, thus facilitating tier-to-tier storage migrations and eventually purging of the data when it exceeds its legal retention period
- Single point of access, thus simplifying the image enabling of the EMR system—the key to achieving Meaningful Use of the images

Arguably, the VNA is the proper approach for enterprise medical image data management. Unfortunately, the properly configured VNA—a dual-sited, mirrored configuration—is an expensive solution that adds an approximately 3X multiple to the cost of a simple mirrored storage solution. That additional 3X cost is associated with the following:

- The VNA software license
- The hosting servers
- The professional services associated with:
 - Deployment of the VNA
 - DICOM data migration from multiple PACS to the VNA
 - Adaptations of individual PACS to support working with a foreign archive

A review of a specific example based upon data from a project² recently completed illustrates the following: the five-year total cost of ownership (TCO) for a dual-sited, mirrored storage solution sized for a community hospital performing 200,000 radiology procedures per year—and already maintaining 82 terabytes of uncompressed historical data—would be approximately \$720,000. For the same hospital, a five-year TCO for a properly configured, dual-sited VNA, using the same storage solution, would be approximately \$2,642,000. That \$1.9 million cost differential is 2.7 times the cost of the consolidated storage solution alone. For numerous healthcare organizations, that 3X differential is difficult to cover in a single budget year, no matter how many critical issues the VNA would resolved.

Phased evolution toward a VNA

An alternative to deploying the full VNA in a single phase is a multi-phase strategy, which coincidentally is the same strategy employed in the early years of radiology PACS deployments. The term “phase” is used to describe a strategic step in a large project or system deployment; when the deployment spans a single year, a series of consecutive phases comprise the implementation schedule for the entire system. The deployment of large projects or systems may actually span multiple years or budget cycles, in which case each phase represents the deployment of a subset of the overall system and subsequent phases build upon the subset deployed in the previous phase. Here, the term “phase” refers to the latter meaning, breaking up the large VNA solution into multiple phases that match yearly budgets.

The first phase of a multiphase VNA deployment strategy often involves addressing one or more significant challenges in the organization with a subset of the VNA configuration that fits within the first year budget.

An increasingly popular first phase of a VNA project is frequently referred to as a Pro-Active Data Migration, which means moving the organization’s DICOM Image Data being managed by one or more department PACS through a basic VNA appliance to an independent storage solution. Pro-Active Data Migration of image data before the individual department PACS are actually replaced can shave considerable time and cost from the inevitable migrations that will occur when the organization chooses to replace any of its existing PACS. During this migration process, the VNA appliance “normalizes” all of the proprietary metadata elements that have been systematically introduced into the image headers by the various PACS, resulting in a vendor-neutral image database that is ready for use by whatever future PACS may be chosen.

The cost of the Pro-Active Data Migration phase includes the following: the DICOM data migration services, the basic VNA application license, and the cost of the independent storage solution. Even though this package should comprise significantly less than 50 percent of the cost of the full VNA project, the price may still be beyond the first year budget. An even more compact (but useful) first phase strategy is to simply deploy an independent storage solution onsite and use a media migration process to consolidate the image data from each department PACS into it—a relatively easier option, because the data format is not modified in any way.

² Private Assessment Project by Michael Gray, May 2011

The ideal time to deploy an independent storage solution is when the existing PACS storage is nearing its third anniversary. Maintenance costs on storage solutions rise significantly beyond the third year, which is generally when the storage volume typically sold with the PACS nears capacity. Furthermore, adding similar storage to a three-year-old storage solution is most likely an investment in old technology.

Many storage solutions exist for this simple independent storage solution phase, but the list will be limited to storage solutions approved by the individual PACS vendors for use with their specific PACS. The storage solutions on the approved list should then be carefully investigated for the following attributes:

- Compatibility with VNA architecture
 - support for multiple interface options
- Support for VNA functionality
 - data duplication
 - automated tier-to-tier media migrations
 - ability to accept and store non-DICOM and non-image data, as well as standard IT infrastructure needs
- Fully scalable and highly efficient

The EMC Isilon scale-out NAS solution fits this criteria and is worthy of closer examination.

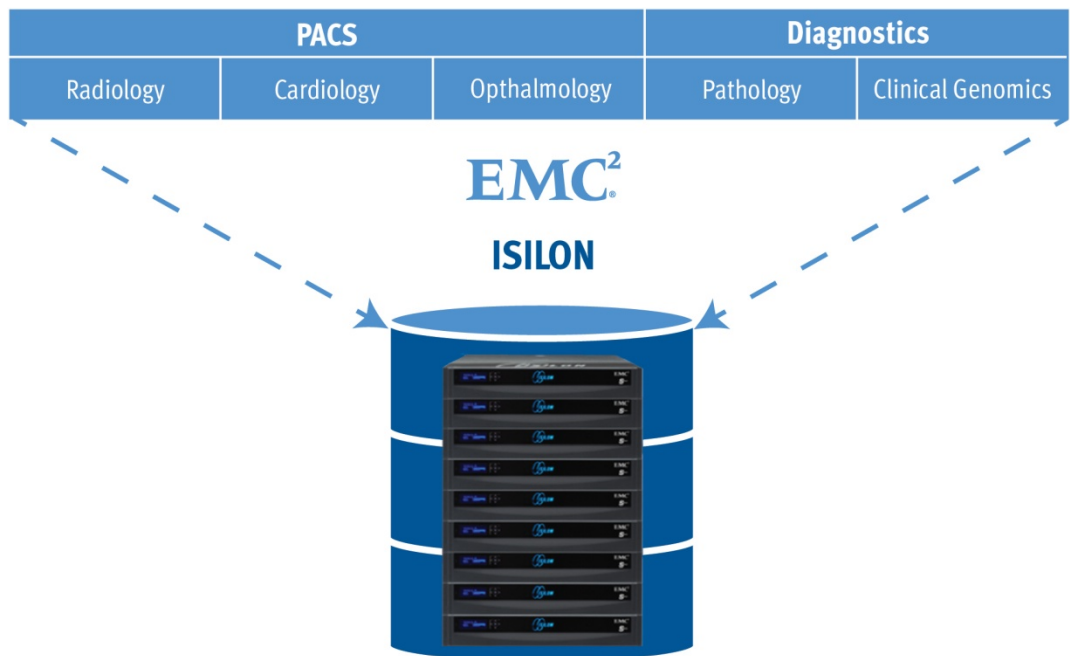


Figure 2. EMC Isilon as a phased approach to VNA

The bridge to VNA: scale-out storage

The EMC Isilon new generation storage solution manages data through a single file system namespace rather than at the block level, and storage appliance nodes are arranged in clusters that support massive scalability. A short description of this storage solution with its features listed below will confirm its suitability as both an independent storage solution for multiple department PACS and as a storage solution for the full featured VNA.

Simple

EMC Isilon OneFS[®] operating system combines the three layers of traditional storage architectures—the file system, volume manager, and RAID (Redundant Array of Inexpensive Disks)—into one unified software layer, creating a single intelligent distributed file system that runs on an Isilon storage cluster.



Figure 3. OneFS eliminates need for separate file system components

EMC Isilon scale-out NAS hardware provides the appliance on which OneFS distributed file system resides. A single EMC Isilon cluster consists of multiple storage nodes, which are rack-mountable enterprise appliances containing memory, CPU, networking, NVRAM, storage media, and the InfiniBand backend network that connects the nodes together. Hardware components are best of breed and include the benefits from the ever-improving cost and efficiency curves of standardized hardware. OneFS allows nodes to be incorporated or removed from the cluster at will and at any time, abstracting the data and applications away from the hardware.

Scalable

EMC Isilon provides a high-performance, fully symmetrical, cluster-based distributed storage platform and includes linear scalability with increasing capacity—from 18 TB to 16 PB in a single file system—as compared to traditional storage.

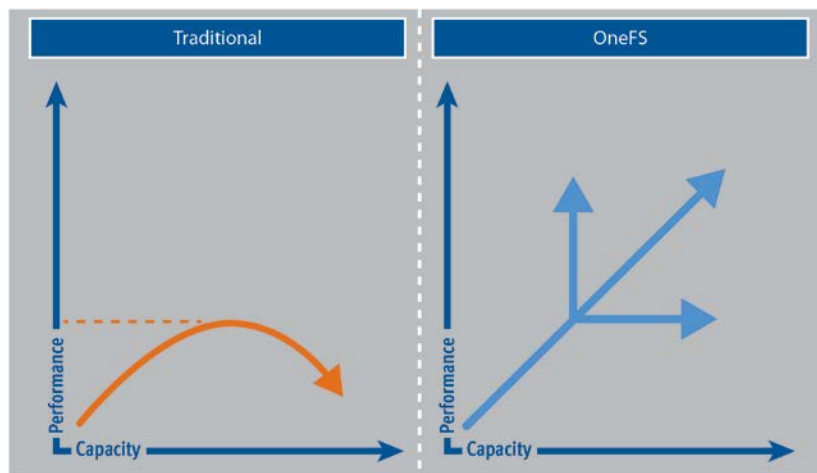


Figure 4. Linear scalability with OneFS

Predictable

Dynamic content balancing is performed as nodes are added or data capacity changes. This makes the storage scale transparently, on the fly, from 18 TB up to 16 PB, without added management time for the administrator or increased complexity within the storage system. The EMC Isilon storage reporting application, InsightIQ, can be used to plan the growth of a system from storage statistics— both for infrastructure and for budgeting.

Efficient

Compared to most storage platforms that use RAID methodology with average efficiencies of 50 to 55 percent, OneFS provides over 80 percent efficiency with its utilization of raw storage³—independent of the location of CPU or compute or cache. This efficiency is at the application level and tiered by the performance types:

- S-Series node for high performance (I/Ops)
- X-Series node for high throughput
- NL-Series node for archive
- NL-Series node for archive

³ See EMC Isilon 80% utilization capability: <http://www.isilon.com/press-release/isilon-guarantees-value-and-simplicity-scale-out-nas>




Figure 5. Storage tiering based on node type

Data is automatically reorganized to optimize performance or capacity. The tiers in the storage cluster are identified as “pools” and managed by the EMC Isilon SmartPools[®] application. A pool is a group of similar nodes that is defined by the user and is based on the functionality or workflow.

A pool is governed by policies which can be changed based on needs; default policies are built in. Policies can be defined by any standard file metadata: file type, size, name, location, owner, age, last accessed, etc. Data can be migrated from pool to pool. The timing for this data movement is configurable: default is one time per day at 10:00pm.

Available

Flexible data protection occurs during power loss, node or disk failures, loss of quorum, and storage rebuild. OneFS avoids the use of hot spare drives and simply borrows from the available free space in the system in order to recover from failures. This technique is called “virtual hot spare.”

Since all data, metadata, and parity information is distributed across the nodes of the cluster, the Isilon cluster does not require a dedicated parity node or drive, or a dedicated device or set of devices to manage metadata. No single node can become a single point of failure. As a result of this feature, the cluster is self-healing.

Enterprise-ready

Snapshots, replication WORM, and quotas are accessible through a simple Web-based UI. Connectivity is supported through standard file protocols: CIFS, SMB, NFS, FTP/HTTP, iSCSI, and HDFS.

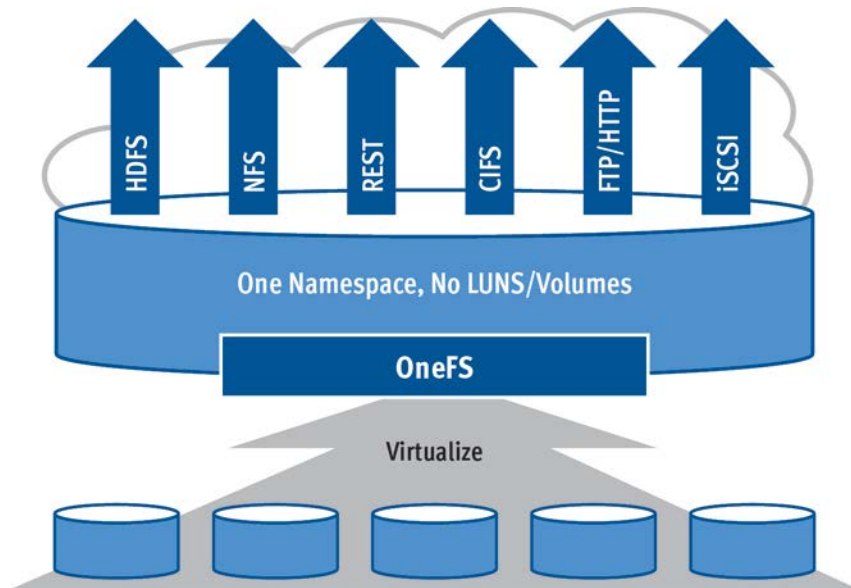


Figure 6. Standard protocols via OneFS

Data is given infinite longevity, which future-proofs the organization from evolving hardware generations and eliminates the cost and difficulty of media migrations and hardware refreshes. Standardized authentication and access control are available at scale through Active Directory (AD), LDAP, NIS and local users. Simultaneous or rolling upgrades to OneFS are possible, with little or no impact to the production environment, and OneFS management software is automated to eliminate complexity.

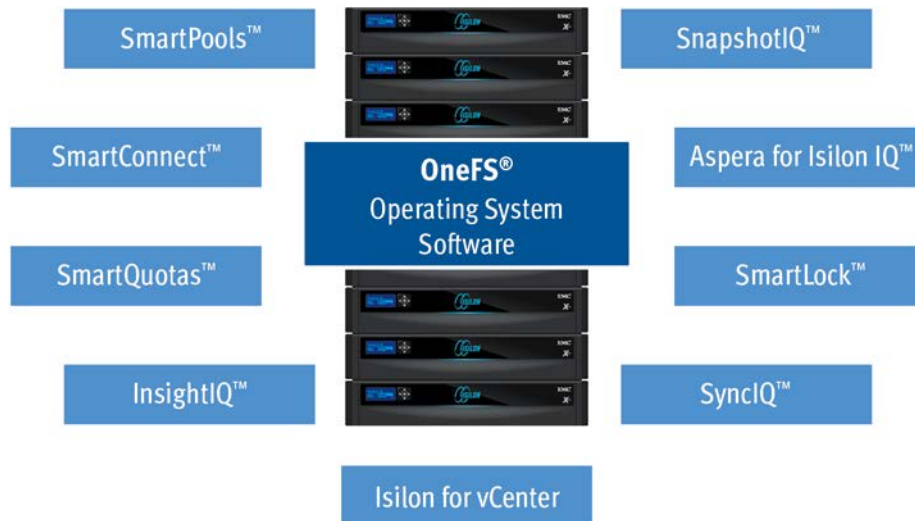


Figure 7. Software suite to manage storage

All of the applications shown above are available as software licenses and are Web-based through the main administrative user interface. A comprehensive command-line based administration is also available.

Connectivity

As previously mentioned (and illustrated in the following graphic), a number of options exist for interface to the department PACS, the future VNA, and any non-DICOM or non-image data sources that the organization might wish to store in the enterprise archive.

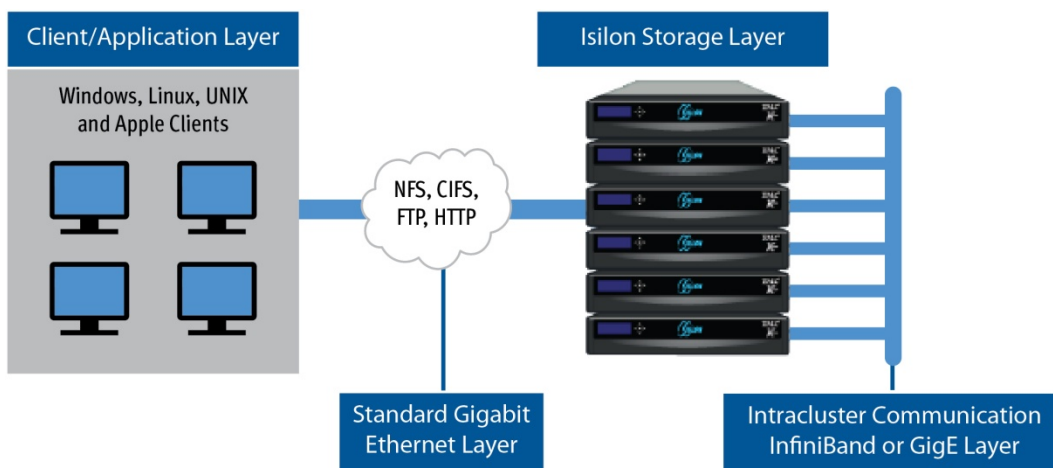


Figure 8. Ease and flexibility of connecting to EMC Isilon storage

Security and data availability

EMC Isilon and Vormetric have teamed up to deliver a highly scalable clustered NAS solution, along with encryption and key management capabilities. This solution protects sensitive information and facilitates compliance with a variety of security requirements. An alliance with Varonis provides log and audit information for regulatory compliance.

SyncIQ replication software provides the data-intensive healthcare organization with an unparalleled solution for replication over the wide area network (WAN) and local area network (LAN) for disaster recovery, business continuity, disk-to-disk backup, and remote disk archiving.

Your partner in this evolution: EMC Isilon

EMC Isilon scale-out NAS key features—especially those similar to the features and value propositions of the VNA—make it an ideal independent storage solution that can be shared by the individual healthcare departmental PACS.

The first key feature is the EMC Isilon single file system (as opposed to block storage). Similar to the Information Lifecycle Management functionality in the VNA, the file location never changes, even when data is migrated from high performance to less expensive storage nodes based on user-defined rules. Simplified data migration between old and new media also eliminates costly data migration and manual data movement between multiple price/performance tiers. The single file system also results in better storage utilization rates and a much better cost ratio to useable space.

A second key feature is the ability to create/assign storage pools, which can be dedicated to different node storage hardware like SSD, SAS, and SATA, or within each storage type, but beyond the scale of a Storage Area Network (SAN). In this case, each department PACS can be assigned its own storage pools. The user can create policies to move image files among different price/performance tiers of storage based on their clinical relevance—for example, age of the study, relevance as measured by last time accessed, or last time changed. Several default policy templates are available, which the user can use as-is or modify. The Vormetric security application available for EMC Isilon storage enables each entity to encrypt data with its own encryption key, so adjacent entities or administrators cannot view encrypted data. This is similar to the VNA ability to securely manage image data submitted by separate departments PACS or various facilities in separate partitions.

A third key feature is data duplication, which allows the organization to automatically create a second disaster recovery copy of all of its image data in a mirrored storage solution located in the organization's second data center. This cost-effective data protection solution supports VNA functionality.

Data Security in the form of built-in data encryption is the fourth key feature of the EMC Isilon solution. A report⁴ released by ID Experts and Ponemon Institute found that data breaches in U.S. healthcare organizations have cost more than \$6 billion a year. According to Mahmood Sher-Jan, vice president of product management at ID Experts, “an incident may not become a breach if it’s encrypted properly.” The ability to encrypt data both in flight and at rest is a valuable prevention measure.

Lastly, there is the rather interesting capability of the EMC Isilon storage solution to manage the metadata associated with the image data in its own storage pool. At some point in the near future, one would expect that Isilon would enable an external application—something other than the department PACS that is using the storage solution for its long-term archive—to independently access the image data using this identifying metadata. Many useful applications like VNA and zero-client, server-side rendering clinical viewers could then use the Web services (HTTP/REST API) interface to directly access the image data.

Conclusion

The most obvious argument for adoption of the VNA is the consolidation of organizational image data in a single repository. Managing all the image and non-image data in the patient’s longitudinal medical record in a single consolidated data repository is less expensive and more efficient than the practice of managing all the data in individual departmental PACS and information systems. Since a dual-sited, mirrored VNA solution may be too costly for a healthcare organization to fund in a single phase, an approach reminiscent of the early years of radiology PACS deployments is recommended: spread the VNA deployment over multiple phases and over several budget cycles.

A smart and comparatively affordable first phase strategy in this approach replaces the long-term storage solutions from those disparate departmental PACS with a consolidated independent storage solution shared by all the PACS. Even if the image data remains in the originating PACS format, the consolidation alone can reduce capital and operational expense, and the technology upgrade can extend beyond a mere improvement in data storage.

The EMC Isilon scale-out NAS solution offers more than a storage upgrade to departmental PACS. With key features such as a single file system, storage pools, data duplication, security, and metadata management, the EMC Isilon solution directly aligns with the data management requirements of the VNA. It is a simple, highly scalable, predictable, efficient, available, and enterprise-ready solution, with advanced features that offer a wide range of connectivity options for both current PACS and future VNA environments.

The EMC Isilon scale-out NAS is a smart and affordable first phase deployment of a multiphase VNA deployment strategy that can easily be built upon through successive phases until a fully featured, dual-sited VNA configuration is realized.

⁴ Benchmark Study on Patient Privacy and Data Security; Sponsored by ID Experts, Independently conducted by Ponemon Institute LLC, November 9, 2010