



Distributed Disaster Recovery and Archival Platform for Medical Systems

Healthcare organizations invariably have numerous modalities, imaging systems, and other medical applications distributed throughout clinical departments without common access to IT infrastructure or best practices. These systems located in the distinct “ology” departments are often acquired and administered by departmental staff members each running their own procedures for protecting data and preparing for system recovery in the event of a failure. While these systems are often not as prominent from an IT perspective as applications like Radiology and Cardiology PACS, they are nonetheless evolving in the same direction with the same important consequences for the business.

Problem: These systems and procedures generate increasingly massive quantities of patient-related data that require long-term and compliant retention. They also increasingly need to be highly available for clinical access with redundancy and rapid failover or recovery capabilities. Usually, the application vendors provide clinical healthcare users with procedures to routinely backup or replicate both the data the applications create and the data and system state required to rebuild a system in the event of failure. However, the lack of a common storage and network infrastructure for supporting these procedures leads both to excessive and redundant costs and limited reliability. The problem is amplified in multi-site healthcare networks where there can be redundant costs for equipment, procedures, and resources across dozens of facilities.

Solution: While every application has its own data structures and vendor-specific procedures for protecting and rebuilding data and systems, all these applications have one thing in common: they need secondary storage repositories where backups, replicas, and archived data can be safely routed and kept available. Healthcare organizations can greatly reduce the aggregate cost of this infrastructure by providing it as a central service to the “ologies” rather than leaving the departments to provision each application environment in isolation. Consolidating and standardizing the secondary storage platform also can greatly improve the organization’s ability to respond to failures of all sizes.

This paper describes how healthcare networks have deployed a distributed and standardized secondary storage infrastructure from ProStor Systems at minimal cost while greatly enhancing the ability of clinical departments across multiple sites to recover failed systems, protect and retain data for long time periods, and resume operations at a secondary location in response to disaster.

Secondary Storage Solution Components

A secondary storage solution capable of addressing a range of applications is not particularly complicated. There are three primary layers of functionality that should be considered:

- A standardized common interface for data protection applications and procedures to submit and retrieve data without modification to normal operation
- An object management layer for abstracting data from physical storage, packaging it for secure storage, and creating and distributing multiple copies to support different recovery scenarios
- The physical storage layer where long-term retention functionality must be matched with minimal cost structure

- » Content user audit trail reporting
- » Automatic single-instance file feature
- » Removable RDX disk based media for offsite vaulting, preservation and disaster recovery
- » Content Addressable Storage (CAS) with sophisticated hash algorithms is used to track and verify all content – ensures the immutability of the data.
- » Can be configured to make copies of data automatically
- » Powerful AES-256 encryption for RDX removable cartridges – provides data security.
- » Encryption keys are managed internally simplifies management of very secure data.
- » Media is backward/forward compatible
- » Reduces backup costs by 50%, cooling and power saves 75%
- » Improved operational efficiencies
- » Ease of use set and forget GUI
- » Data compression available
- » On-the-fly file level data de-duplication
- » 1TB/cartridge

Common Standard Interface

Over the last several years, NAS file systems have become the interface of choice for allowing most content management applications to move data of all types into secondary storage. Historically, magnetic disk-based storage was not economically viable for storing inactive copies of data outside of the production environment. Application vendors would therefore have to develop a means of interfacing with and testing and validating secondary storage technologies such as magnetic tape and magneto-optical disk. However, magnetic disk-based systems are now in fact more affordable than tape and optical and this has greatly simplified the interfacing of applications to secondary storage.

Virtually all medical content management and HIS applications support writing secondary data copies – including backup savesets, snapshots, and archives – as files to a shared disk volume mounted on the network via CIFS or NFS. Secondary storage systems require a wide range of functionality to properly process data and ensure that it is stored compliantly and efficiently in multiple copies at multiple locations. However, from an interface perspective they should simply present a NAS file share that can be mounted by the application across the network, written to via standard network protocols, and supporting native network security. In addition to NAS, medical imaging applications almost universally support the DICOM protocol for transferring medical imaging studies to external environments. As a result, a general purpose medical archive repository may typically also require a DICOM interface providing users and applications full search and retrieval capabilities based on the critical meta-data stored in DICOM tags.

Object Management Layer

The principle goal of secondary storage is to ensure that data is secure and will not be lost, and that particular system states can be reconstructed on demand even through a range of failure scenarios and over long time periods. To make this possible, the secondary storage environment has come to be viewed as more than a single physical storage location where data is deposited. Secondary storage has, in fact, now come to be viewed as a self-managing automated system, or “object store,” leveraging the network to operate across multiple physical locations. Secondary storage appliances therefore have their own object management capabilities, making multiple copies of data and distributing them in different locations so that no particular event could result in total data loss.

- If production data is lost during normal application operations, the data should be easily recoverable from a local secondary data source within the secondary storage system.
- If a single application/system fails, a local data copy is the most convenient source for rebuilding.

- In a site failure, systems may have to be rebuilt at a secondary location and it is a requirement that the secondary storage system span multiple locations.

When you factor in long-term and compliant retention needs, there are numerous other benefits to building a data management capability into the secondary storage. Increasingly mandatory, functions are made possible through an object management layer within the secondary storage system.

- Data can be assigned a unique digital “fingerprint” that allows for subsequent retrievals to be authenticated, guaranteeing data integrity on retrievals that might be performed years after the data was originally ingested into the system.
- Data that is stored for the long-term, especially on removable media, should be encrypted to comply with security standards.
- Data can be compressed and deduplicated to reduce physical storage costs.

Secondary Physical Storage Considerations

The amount of data that needs to be preserved on secondary storage for subsequent access as either a long-term archive or as part of disaster recovery procedures is growing profoundly. According to IDC, the multiple copies of this data required to ensure data is causing data storage needs to grow at a CAGR of 60%. Therefore, the physical secondary storage medium needs to be cost-effective, scalable and enduring.

Physical secondary storage needs to not only have a relatively low cost of acquisition, but also must minimize other cost factors, such as operational overhead and power and space consumption. Since data on secondary storage is retained for long periods of time, the cost of future growth and upgrades need to be factored into the total cost of ownership.

Another consideration is the type of media used in the secondary system. There are considerable advantages to using removable or portable media.

- Removability can ensure that at least one secondary copy of the data is physically safe by virtue of being offsite, preferably in a “vault.”
- Removable media often provides an important alternative to WAN-based data transfer when organizations find they have limited bandwidth.

In terms of total cost of ownership (TCO), magnetic disk has recently become the favored medium for physical secondary storage over the traditional low-cost, tape option. While still more expensive than tape on a \$/GB basis, magnetic disk avoids the high overhead costs that tape devices and robotic libraries introduce to the total acquisition cost of tape systems. However, tape is still often favored because of its removability and the fact that when it is not in use, it does not draw power.

Recently, RDX – a spin-down removable disk media alternative -- has emerged as an ideal secondary storage media building block.

RDX provides both a magnetic disk-based format and a tape-like, cartridge-based packaging that allows data to be easily and securely removed to offsite locations which greatly reduces organizations' power consumption for inactive data.

ProStor Systems RDX-Based Secondary Storage Solutions

ProStor Systems is a leader in providing cost-effective, long-term retention of digital information. ProStor is recognized for developing and promoting the widespread adoption of its RDX® removable magnetic disk technology. The company also develops and sells the ProStor InfiniVault® data storage system -- an automated and comprehensive secondary storage platform that leverages RDX to address the expense and complexity of protecting long-term, regulated data, making it easily available and ensuring compliant data management throughout its lifespan.

ProStor's RDX removable magnetic disk technology is designed to provide a unique level of protection and economical storage for valuable information over an extended period of time. An RDX-based system consists of one or more rugged, removable disk cartridges and either an individual dock or multi-cartridge enclosure, the latter associated with the ProStor InfiniVault secondary storage system. When the RDX cartridge is inserted into the dock or enclosure, it functions as a normal disk drive. Yet it can also be used and removed for off-line storage just like a traditional tape or optical drive with backups or for archival data. The media, which includes a 2.5-inch mobile disk drive with ramp-up heads housed in a shock-proof casing, can be safely archived for 30 years¹. RDX disks provide data transfer rates of over 300GB per hour and file retrievals within milliseconds. RDX provides an unusually strong economic model.. Not only does it leverage the constantly decreasing commodity pricing of SATA, but it is modular, allowing organizations to acquire additional capacity after newer generations with superior price, density, and performance attributes become available. Finally, unlike fixed array-based disk, it does not require annual maintenance fees, even though it delivers a 30-year life-span.

InfiniVault - A General-Purpose Secondary Storage Platform for Healthcare

InfiniVault is an intelligent magnetic disk-based storage platform that incorporates a powerful object management capability and leverages the unmatched functionality and economics of RDX. InfiniVault provides

both NAS and DICOM interfacing so that a range of applications in the healthcare environment can easily and simultaneously exploit its virtualized storage capacity to meet their unique archival and disaster recovery requirements. Thanks to its object management layer, the InfiniVault is able to present infinite virtual storage "vaults" to any number of distinct applications or data classes within a health network location. Each virtual storage vault is set up independently with automated rules that determine how long data will be retained, the protection methods to be used (such as WORM and encryption), and the number of copies and their disposition on physical media across the health organization's distributed environment.

Performance

InfiniVault brings all the benefits of magnetic disk performance to secondary storage while preserving the removability, portability, and durability of traditional offline or near-line media. Unlike tape or optical systems, InfiniVault supports removable media without robotics, avoiding the access latency found with tape or optical from mounting and positioning the media. To be online and accessible, RDX cartridges simply need to be connected through the InfiniVault's standard SATA electrical interface. InfiniVault systems provide a modular array of switched slots that allows all connected RDX cartridges to be on-line simultaneously and, when active, accessible with standard magnetic disk access speeds. InfiniVault systems also provide fixed disk RAID cache to ensure that the system can support typical NAS data transfer speeds. The cache is sized for each InfiniVault model with sufficient capacity to make recently created data available at primary storage speeds. The InfiniVault Model 100 comes with 13 TB of cache front-ending 20-100 RDX slots (Currently 20-100TB of data).

Data Protection and High Availability

The InfiniVault properly protects and can ensure high availability for archival and disaster recovery healthcare data. It automatically writes multiple copies of data to RDX media cartridges that can be locally connected and accessed online, removed from the system and taken offsite, or replicated to other remote InfiniVault systems via the WAN. In a typical scenario, an InfiniVault might be set to keep one RDX copy of data online in the system with a second copy removed to a vault to prevent loss of data in all failure scenarios. Another copy of the data could be replicated to a secondary site providing the means to quickly and easily re-direct applications to an alternative network location for continuous data access during a primary site disaster. After failover has occurred, rebuilding or replacing an InfiniVault and reverting operations is easy and rapid. The InfiniVault routinely and automatically backs its own metadata up to both RDX cartridges and to remote network locations. This database

Compliance

Depending on internal policies and location, different healthcare organizations have different requirements for how patient-related data is retained. The InfiniVault and RDX were designed to ensure compliance with even the most stringent policies. First, the InfiniVault allows different data classes to be assigned to different virtual vaults each with their own granular rules for managing retention, redundancy, and security. As described above, the InfiniVault automatically ensures fundamental protection against data loss through multiple copies at different secure locations. It also uses digital fingerprinting to ensure data integrity and authenticity. Each class and sub-class of data is automatically managed in terms of retention and the system can be set to leverage hardware and software-mediated WORM controls, preventing data from being overwritten or modified in any way prior to its expiration. Since RDX provides a 30-year life-span, most health data can actually reside continuously on an RDX cartridge without the need to be regularly transferred to new media as part of a storage "refresh."

From a security perspective, the InfiniVault integrates with the native network and application security and access control schemes. Internally, the system has its own highly controlled roles-based access controls, preventing access to underlying patient data even for administrators. The system provides multiple levels of security to prevent unauthorized access to data stored on RDX. First, firmware in the InfiniVault prevents RDX cartridges from being read on any system other than the InfiniVault where it was created. Second, the InfiniVault has physically secure locked enclosures that prevent online RDX cartridges from being removed from the system. Third, InfiniVault uses advanced AES256-bit encryption algorithms to encrypt the data at rest on RDX and in transit during replication. Even if stolen, removed RDX cartridges will never provide access to health data, no matter how sophisticated the individual attempting access. Finally, compliance includes the ability to audit data access attempts and report on the chain of custody of data. InfiniVault explicitly tracks the chain of custody of data and provides standard reports to demonstrate that data has been compliantly handled and identify when attempts to compromise data have occurred.

Unmatched Economics

While providing important functionality, secondary storage nonetheless holds data outside of live production and is expected to be significantly less costly than tier-1 capacity. The InfiniVault provides a unique set of economic advantages based on its packaging and leveraging of RDX. First, RDX uses standard SATA disk components precisely to exploit the significant cost reduction trajectory of commodity disk products like SATA. Furthermore, the modular packaging of RDX allows organizations to invest in an InfiniVault only buying the immediate RDX capacity needed.

As more capacity is required, they can then take advantage of the higher capacities and lower price per gigabyte of newer RDX generations. For instance, 2010's 1TB cartridges will be available at almost the same cost in 2011 in a 1.5 TB configuration. Fortunately, the InfiniVault can mix and match different generations of RDX without any constraints. Another major advantage of RDX is that it does not involve the annual maintenance and support costs that users experience with fixed disk storage products. And the RDX format is unique in that it can be used for up to 30 years. Unlike standard disk storage arrays that must be refreshed every 3-5 years, RDX capacity can be kept in production for up to 10 times as long, averting significant repeated capital outlays. With InfiniVault virtualizing the underlying storage through its object management layer, there is also a benefit that long-term data need not be migrated from old to new storage with an operational impact on applications. Instead, the InfiniVault automates the process of intelligently re-distributing data – if desired -- transparently to the applications above it. For all these reasons, InfiniVault solutions are significantly more affordable than most other secondary storage offerings.

Sample Large-Scale Health Network Deployment

The diagram below shows a real-world healthcare network's enterprise-wide deployment of multiple InfiniVaults to provide a unified secondary storage infrastructure for DR procedures for over 75 miscellaneous modalities, imaging and HIS systems. The organization chose InfiniVault and RDX to meet this need and elected to deploy various InfiniVault models at each of their 9 regional hospitals, sized according to the local data requirements. The systems appear on the network as CIFS- mounted file shares. There is also a DICOM data transport interface. Each of these InfiniVaults provides virtualized storage capacity to support a local on-line copy of archive, backup, and DR data on RDX media maintained in the system. Different clinical and HIS applications each have their unique procedures and tools for creating secondary data copies. For instance, imaging systems perform system state backups, database dumps, and continual image archiving. All of these processes write to specific sub-folders in appropriate virtual vaults that then automatically apply the appropriate data management rules determining retention times, security, and other relevant policies. Data written to these systems lands on the fixed-disk cache and is available for recall over some time frame directly from the cache. Immediately upon ingestion a primary copy is made to RDX volumes maintained in the InfiniVault system. Secondary copies are made to additional RDX cartridges that are then removed to a local vault for protection against local disaster. Finally, the data is also replicated to large remote InfiniVault systems at one of two central data centers which are well positioned and have network bandwidth sufficient to recovery and support application remotely on behalf of the smaller hospitals.

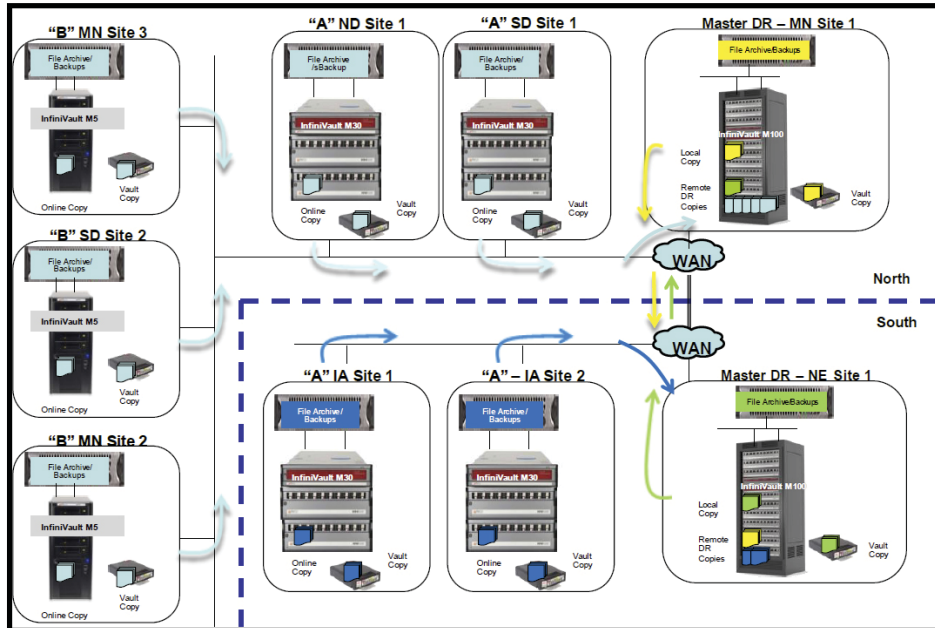


Figure: Real World InfiniVault Application Showing Network and Data Copy Distribution

This network of InfiniVaults, spanning all nine of the major hospitals in the enterprise, ensures that when Healthcare applications need to be recovered, they benefit from a standardized, reliable, and well-managed storage environment. The 3-copy data strategy properly distributes data so that individual system rebuilds can be performed locally and quickly, while all systems can potentially be rebuilt at a remote site in the event of a primary site failure. Removable copies also ensure that local operations can be rebuilt after a disaster without a massive data transfer across limited WAN resources.

Summary

A clear need exists for many healthcare organizations to provide their clinical departments with an efficient and easy-to-access infrastructure capable of meeting the secondary storage requirements for data protection, archiving and disaster recovery of a wide range of distributed clinical applications. To properly meet this need, a secondary storage system should provide several layers of functionality including: a universal interface, a robust object management capability, and storage media optimized in terms of price and functionality for long-term data retention. The InfiniVault secondary storage system and its component RDX removable disk technology provides these layers and can be deployed in a cost-effective manner across a distributed healthcare organization to provide a powerful infrastructure solution for medical system disaster recovery and long-term healthcare data retention.



/ An Avnet Company /

» www.rorke.com