IMPACT OF HIPAA ON HEALTHCARE DATA MANAGEMENT

WHITE PAPER

DON BAUNE
JOHN GALASSO

StorageTek®

TABLE OF CONTENTS

GLOSSARY

ASP — Application Service Provider

CR — Computed Radiology

CT — Computed Tomography (CAT Scan)

DICOM — Digital Imaging and Communications in Medicine

HIPAA — Health Insurance Portability and Accountability Act

HSM — Hierarchical Storage Management

ISS — Internet Storage Supplier

JCAHO — Joint Commission on Accreditation of Hospital Organizations

MRI — Magnetic Resonance Imaging (MR Scan)

PACS — Picture Archive and Communication System

SSP — Storage Service Provider

TESS — Transactional Enterprise Storage Solution

INTRODUCTION

Healthcare enterprises have three data management/storage needs: backup and restore, disaster recovery and patient image archiving. HIPAA has added the significant additional burden of privacy and security of protected health information. The purpose of this document is to describe a system that supports good data management practices for healthcare in the HIPAA era.

As the healthcare enterprise begins to develop strategies that will conform to federal laws, oversight agencies and accreditation reviews (HIPAA, FDA, JCAHO), it is important to take an enterprise view of data management, patient archival, disaster recovery and data security needs. As the need for a central data policy is mandated, the days of being able to ignore the consequences of "information islands" in the healthcare enterprise are quickly coming to an end. A coherent enterprise data management strategy and a centralized data repository lead to a total cost of ownership that is less than the traditional "grow as needed" departmental storage solution.

In the past, installation of a centralized data archive solution for a healthcare enterprise has required a certain amount of guesswork to determine the final hardware and software configuration. Now that several successful installations are in operation, much of this guesswork has been eliminated. Proven functioning systems serve as models.

INFORMATION ISLANDS IN THE HEALTHCARE ENVIRONMENT

The traditional role of the IT department in a healthcare enterprise was focused primarily on administration, i.e., patient registration, billing, insurance and possibly medical records. The amount of data generated by these activities was small when compared to the total data generated in the healthcare enterprise.

As digital modalities (CT, MR, CR, digital mammography, and digital cardiology) were introduced, individual departments wanted to store the raw digital data as well as a representative subset of this data in the form of films, which were the legal standard. As digital reading became the norm for interpreting and diagnosing pathology, storing these images in a film format was no longer clinically effective or financially prudent. A typical spiral CT scan can contain over 2,000 images, requiring over 100 sheets of film.

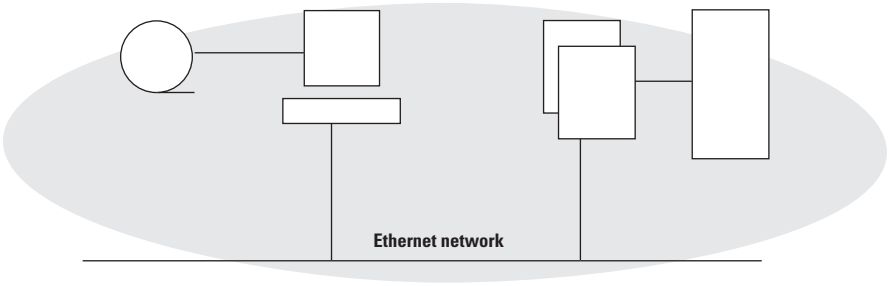

Typically a SCSI attached archive was connected directly to the PACS database server. Data was stored in the archives in a closed proprietary format and could only be accessed through the PACS access interface — a structure that applies to the stand-alone tape and optical drives provided with MR and CT scanners today.

Network infrastructure was not capable of handling the amount of image data that digital modalities generated. Furthermore, this data was viewed as clinical data outside the administrative responsibility of IT. Departments began making their own digital archives. IT was often unaware these archives were being installed outside its control and influence.

It is now common to find multiple storage strategies throughout the enterprise resulting in departmental islands of patient data. Unfortunately, most of these archives are under the control of departments with little or no IT expertise and without strict IT management controls and safeguards.

As data needs have grown and storage technology has advanced, archives in radiology and cardiology departments are typically stored on multiple legacy systems. Ultrasound, nuclear medicine and digital mammography also typically maintain dedicated archives. Specialized databases include neo-natal growth charts and risk classifications, pharmacy databases, pathology databases and labs.

Each of these servers represents a data island and a potential point of privacy breach of patient data that could be released to the public. Each also represents a critical healthcare enterprise activity that needs to have a data policy in place for disaster protection and recovery, access control, and patient care and protection procedures.

IT itself hosts automated backup of the hospital administrative and accounting systems. Non-automated backup of independent departmental servers distributed throughout the healthcare enterprise represent another set of issues in a HIPAA compliant environment. While federal standards (HIPAA) do not specifically state how a hospital will provide data security, information islands make securing health information a daunting task.

APPLICATION STORAGE MANAGER® (ASM) SOFTWARE

Patient data retention requirements, often dictated by federal and state regulations, place a further burden on security. Daily backups can have a life of several days to weeks. Patient medical records, image data and other patient information can have very long retention requirements. Rules vary widely between states but Table 1 shows the potential length of time records may need to be kept.

Patient/proceedure	Potential archive requirement
Adult, age 21	5 – 7 years
Adolescent	5 years
Mammography exam	Life of patient

Table 1

These retention rules often exceed the expected life of the archived media. It is not that the data bits on the media are in danger of becoming corrupted or degraded, rather the technology to read the data becomes obsolete. Today it is difficult to find a working Betamax video tape player, 8" floppy-disk drive or even a more recent 5 1/4" floppy-disk drive. The data on the media may still be viable but the ability to read the data has been lost to technology advances. When the next generation of archive technology is made available, the enterprise then has decisions to make in order to store data for the required time:

- Should we convert the data to the next generation archive technology?
- Is there a way to migrate the data that is not network intensive?
- Will we have to create and maintain incompatible media archives because the data migration from one media format to the new technology is not cost- or time-effective?
- Should we maintain legacy archives through expensive service contracts to ensure parts and service until the data exceeds the legal retention period?
- If we maintain a film copy should we simply abandon the old archive?

Data migration, retention policy, data storage and disaster recovery planning are often not included in the initial design of the PACS archive. These issues can be resolved by using a data broker that separates storage technology from data acquisition and indexing.

StorageTek® Application Storage Manager® (ASM) software is a data broker that has been used successfully in many healthcare enterprises. While not all PACS providers have the ability to use StorageTek ASM software today, it is quickly becoming a healthcare industry standard. ASM software provides the ability to implement IT disciplines in a healthcare enterprise data archive. Although ASM software is not the only product that provides the needed separation between storage technology and data acquisition and indexing, it is currently the most widespread HSM product in use in the healthcare environment. StorageTek ASM software provides the following capabilities.

Data isolation

Data is stored in the most cost-effective form, which is currently tape. In terms of access requirements and cost, ASM software facilitates best-of-breed archiving.

Technology obsolescence protection

Data can be moved as a background function from one archival technology to the next-generation technology. This reduces the load on the main archive server and does not require extensive network resources.

Multiple copies of data

Tape pools are used to generate multiple copies of the data. Tape pools do not have to be identical. A second copy of data can be made to a much higher density tape but with less than optimal access patterns. ASM software can be set up to allow data to be written to multiple tape archives for true continuous disaster protection.

Multiple use archive

The tape library can support many data streams simultaneously. An archive originally designed to serve as a cardiology archive can be expanded to support radiology and IT enterprise backup.

PACS INTERFACE

Until a few years ago, the only methods of image archiving available to PACS vendors were to use one of the few general-purpose HSM products on the market or develop their own HSM. Because of cost and performance, most PACS vendors opted to develop their own. They believed that a specialized HSM tuned to the needs of a PACS archive would be able to out perform a general HSM. For the most part, they were right. However, many PACS vendors failed to recognize that the hardware they chose for data storage continued to change and improve. Keeping the PACS archive current with storage technology took an increasing share of the enterprise's PACS development budget.

At the same time, HSM technology continued to improve to the point that the ASM software product from StorageTek was able to provide performance similar to the dedicated HSMs that had been integrated into PACS systems. ASM had the advantage of being a stand-alone package that continued to improve at a much faster pace than an HSM buried inside a PACS system. ASM gave the healthcare enterprise the freedom to choose the best archive storage solution available when purchasing the PACS system without being locked into choices the PACS HSM supported. This becomes more important with the development of ISSs, SSPs, and ASPs, an increased need for multiple copies in multiple locations and the need for greater availability.

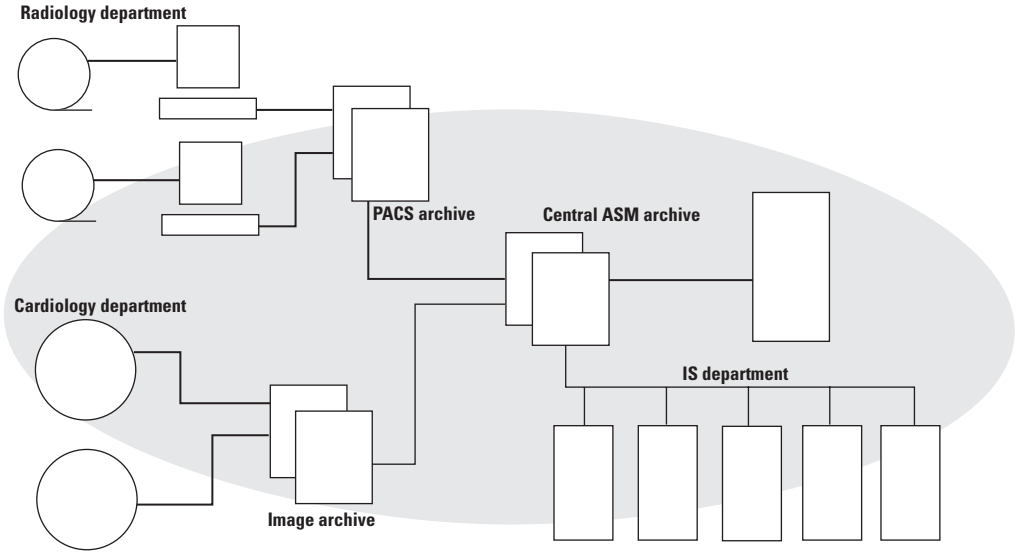

StorageTek's Application Storage Manager® (ASM) software is a self-contained Nearline® archive software system that provides a permanent data storage solution for the PACS archive.

- Data is acquired on any number of digital imaging devices (CT, MR, CR and so on).
- The imaging device sends a copy of the patient exam to the PACS archive using the DICOM communications protocol.
- The PACS archive stores a copy of the image on a local RAID disk and patient demographics on a local database.
- At a predetermined point, a copy of the patient exam is sent to the ASM software server.
- The ASM software server accepts this image and caches it to a local disk drive.
- Using a data policy that is set on the ASM software system, the data is then sent to the tape library or libraries.
- As the local cache on the ASM software server fills to a pre-determined high-water mark, the data is purged from the disk cache. Disk directory entries remain on the server as an index to where this data can be found on tape.

TRANSACTIONAL ENTERPRISE STORAGE SOLUTION

After the image archive is freed from the direct control of the PACS archive, it becomes a sharable resource through the use of ASM software. A number of options become available. The healthcare enterprise can begin to realize the value of having a true Transactional Enterprise Storage Solution (TESS) with benefits that include:

Cost-effective storage

"Giant economy size" truly applies in the digital healthcare enterprise storage model — the larger the archive library the lower the cost per megabyte.

Reduced headcount

A single part-time administrator can monitor and maintain a central storage resource much more easily than several employees managing archives distributed throughout the healthcare enterprise.

Greater security

A single TESS can have additional security procedures and access controls put into place that would otherwise be cost prohibitive if they were distributed throughout the healthcare enterprise. If libraries are distributed (as they traditionally have been), there are many more possible points of security failure.

Data isolation

By using a single archive strategy for the healthcare enterprise, a policy-based storage strategy can be employed and enforced. The healthcare enterprise is able to set policy.

Data access viability

ASM software provides manageable migration paths for the archived data to the next generation storage technology in a controlled, proven process.

Resource utilization

Each department can utilize more resources than it purchased. Let's say IT needs three tape drives to perform enterprise backup in a given time window, radiology needs three drives to perform the pre-fetch and peak daily ad-hoc retrievals, and cardiology needs three tape drives for its peak daily requirements. A tape library with a total of six drives will be able to serve all these needs and give greater throughput and redundancy. Again, if a tape library has three drives, and one was unavailable because of maintenance, one-third of the capacity would be lost and the backups would take thirty-three percent longer. In a larger shared archive the loss of a single tape drive represents a loss of one-sixth of the total capacity and still makes a sufficient number of drives available to get the required tasks completed on time.

17:00 – 22:00	Cardiology batch store	3 drives needed
17:00 – 22:00	Radiology pre-fetch window	3 drives needed
17:00 – 22:00	Enterprise backup	3 drives needed

Table 2

In a traditional captive archive there would be three libraries and a total of nine tape drives. In a shared archive with a total of six drives, three tape drives would be dedicated to the time-based task with three drives available for down drives and emergent requests for archived data from the other services.

Assumptions for multi-departmental TESS

	Radiology	**Cardiology**	**Information systems**
Total work days	254	254	254
Number of procedures	150,000	150,000	
Average size of data set	30 MB	250 MB	500,000 MB backed up/week
Percent of priors	50%	50%	
Number of studies/priors recalled	2.5	2	
HIPAA compliant disaster recovery copy	1 made automatically by ASM software	1 made automatically by ASM software	1 via backup software (VERITAS, Legato, etc.)
Total data MB/day read/write pre-fetch (includes disaster recovery copy)	57,579	44,291	100,000
Total hours for pre-fetch of studies @ 9,000 MB/hr	22,146 MB 1 T9840 = 2.46 hr 2 T9840s = 1.23 hr	14,764 MB 1 T9840 = 1.64 hr 2 T9840s = 0.82 hr	Data backup: 1 T9840 = 11.11 hr 2 T9840s = 5.56 hr 3 T9840s = 3.70 hr
Primary utilization hours (pre-fetch/data backup)	4:00 – 11:00 p.m.	5:00 – 11:00 p.m.	12:00– 5:00 a.m.
Minimum hardware requirement for captive departmental solutions (information islands) is seven StorageTek® T9840 tape drives, three libraries	3 tape drives 2 read /1 write	2 tape drives (avoids single point of failure)	3 tape drives

Table 3

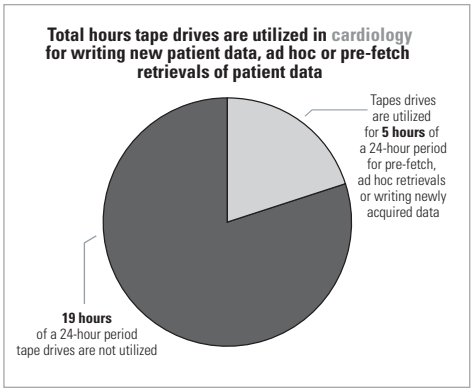

Figure 1

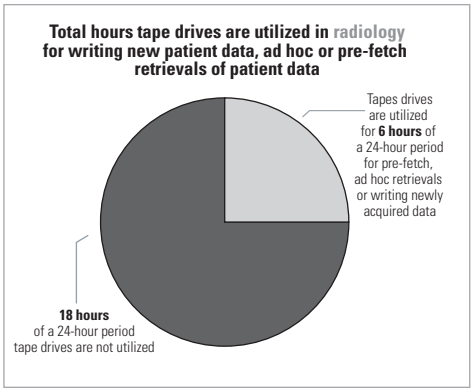

Figure 2

Estimated average performance based on a variety of network environments:

- Tape drive capable of 40 ad hoc retrievals and exchanges per hour
- Sustained transfer rate 10 MB/sec
- Query, load, seek, transfer and unload within 12 seconds
- Transfer rate — 9,000 MB/hr (read/write)

In our example, during peak utilization hours — i.e., pre-fetch or data backup — the ASM software prioritizes the requests in the following order: first priority is read/transfer (ad hoc retrieval); second priority is write new data (backup). The system configuration includes a RAID disk cache because StorageTek® T9840 tape drives can transfer data faster than most networks can, by facilitating data streaming, should multiple ad hoc requests for patient data occur. With a data transfer rate of 10 megabytes per second and an average load, seek, transfer and unload time of 12 seconds, a typical radiology file (100 MB) would transfer in less than 25 seconds and a typical cardiology file (250 MB) in less than 45 seconds. This rapid transfer rate mitigates any impact on the backup window and virtually assures the backup remains unaffected.

Additionally, by pooling the T9840 tape drive resources, the loss of a T9840 tape drive due to maintenance/service issues causes less degradation in system performance. For example, if four of five drives are available for backup and one is lost to a service problem, the backup window remains unaffected. Since repairs are generally accomplished within two to six hours, the department suffers with no loss of performance through resource pooling of the remaining drives.

In a traditional captive archive (information island), there would be three libraries and a total of seven tape drives. In a shared archive with a total of five drives in one library (TESS), four tape drives could be dedicated to the time-based task with one available for down drives and emergent requests for archived data from the other services. In a captive departmental solution, for example, if one of the three recorders suffers from a service issue, the backup of data will not be achieved within the backup window. Sharing the archive facilitates higher performance, mitigates impact on performance due to service problems and lowers the cost for implementation.

After a TESS is implemented, what additional benefits can be gained?

- The healthcare enterprise will have a true inter-departmental resource and no one department has to own the archive.
- Economies of scale are possible that could not be achieved before.
- Standard IT practices can be implemented in the healthcare environment.
- Disaster recovery planning is possible and practical.

COST OF OWNERSHIP OF AN IMAGE ARCHIVE

The total cost of any data archive system includes much more than the initial purchase price — cost that a TESS can help mitigate. Some of the components of ownership include:

Initial equipment purchase

What is the initial cost to purchase the equipment? Short-term reductions in the initial cost can lead to large expenses later in the life of the archive, such as migration costs — an important consideration that must be planned from the beginning.

Maintenance

What is the projected cost for monthly hardware maintenance and downtime for preventive maintenance and unexpected failures?

Media

What is the piece cost per megabyte? As the enterprise grows, can the media be reused or must it be replaced?

Servers

What is the cost to maintain, monitor and connect to the archive servers? Is a single server more cost effective than several? Does a single server serve the long-term needs of the healthcare enterprise?

Data protection

What is the ultimate capacity of the archive? Will several be needed to meet the anticipated growth needs? Will media have to be put on the shelf with an operator taking the role of the archive?

Data obsolescence (migration)

What is the exit strategy? Is there a plan to move the enterprise from the current technology to the next generation? What is involved and has it been tested? Data obsolescence is often the last thing considered when an archive is installed.

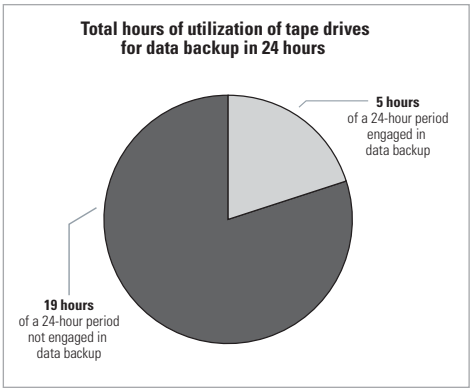

Figure 3

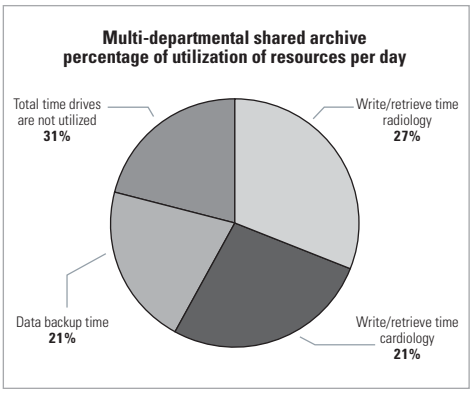

Figure 4

SUMMARY

A TESS managed by StorageTek's ASM software solves data management, security and disaster recovery problems created by information islands in a healthcare enterprise, permits more efficient planning and supports good data management practices for healthcare in the HIPAA era. Periodic fine-tuning of the system prevents performance degradation and facilitates upgrades that prevent data from being compromised.

PACS PARTNERS

- Agfa Healthcare
- Philips Medical Systems
- Rorke Data, Inc.
- Kodak Healthcare Imaging
- Siemens Medical Solutions

World Headquarters
Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA

Phone: 1.800.877.9220

About StorageTek®

Storage Technology Corporation (NYSE: STK), a $2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures, and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.786.7835.

International Offices

Australia: 61.2.9438.4844
Austria: 0800.20.16.31
Belgium: 0800.75.327
Brazil: 55.11.3044.4599
Canada: 905.602.5586
Chile: 56.2.201.0490
China: 86.10.8251.1911
Denmark: 44836111
Finland: 08001.13361
France: 0800.82.83.57
Germany: 0800.181.6238
Hong Kong: 852.8200.0791
India: 91.80.226.7272
Ireland: 1800.55.33.54
Italy: 167.790.852
Japan: 81.3.3746.9711
Korea: 82.2.2191.1100
Malaysia: 603.772.41125
Mexico: 52.55.9177.1800
The Netherlands: 0800.022.8496
Norway: 800.11.220
Poland: 48.22.5757.157
Singapore: 65.6774.9248
South Africa: 0800.99.5820
Spain: 900.99.33.66
Sweden: 020.798.711
Switzerland: 0800.83.87.65
Taiwan: 886.2.6600.9840
United Arab Emirates: 971.4.3900162
United Kingdom: 0800.731.8852

Specifications/features may change without notice.

MZ 9110 C 10/03

StorageTek equipment is manufactured from new parts, or new and used parts. In some cases, StorageTek equipment may not be new and may have been previously installed. Regardless, StorageTek's standard warranty terms (as set forth at www.storagetek.com/warranties/) apply, unless the equipment is specifically identified by StorageTek as "used" or "refurbished." Replacement parts provided under warranty or any service offering may be either new or equivalent-to-new, at StorageTek's option.